KUNGL
TEKNISKA
HÖGSKOLAN

Knowledge Net—A Support for Sharing Knowledge within an Organisation

Kristina Groth

Stockholm 1999

Licentiate thesis
Royal Institute of Technology
Department of Numerical Analysis and Computing Science

KUNGL
TEKNISKA
HÖGSKOLAN

Knowledge Net—A Support for Sharing Knowledge within an Organisation

Kristina Groth

Interaction and Presentation Laboratory
Department of Numerical Analysis and Computing Science
Royal Institute of Technology
SE-100 44 Stockholm
SWEDEN
kicki@nada.kth.se

HMI

NADA

Akademisk avhandling som med tillstånd
av Kungliga Tekniska Högskolan framlägges
till offentlig granskning för avläggande av
teknologie licentiatexamen i Datalogi
torsdagen den 11 mars 1999 kl 13.00 i sal E2,
Lindstedtsvägen 3.
Avhandlingen presenteras på engelska.

ISBN 91-7170-372-1

IPLab-156
TRITA-NA-9902
ISSN 0348-2953
ISRN KTH/NA/R–9902–SE

KTH Högskoletryckeriet, Stockholm 1999

Abstract

Knowledge is today considered the most valuable asset in many organisations. Especially larger organisations may know what competencies they have but not where within the organisation or with whom they can be found. If the knowledge of an organisation can be made accessible within the organisation then the knowledge could be more effectively used. To support the sharing of knowledge many attempts have been made to store individuals' knowledge (often called organisational memory systems).

This thesis focuses on the alternative idea of a so-called *knowledge net*. In a knowledge net information about "who-knows-what", rather than "what" itself, is stored. The focus in a knowledge net lies on individuals' knowledge. A rating method is used to find persons with the "right" knowledge (not necessarily the one who knows most) about a topic. Each "member" chooses what to contribute with to the knowledge net implying a willingness to share the knowledge with others. A knowledge net also includes facilities to contact the suggested "experts".

Three studies have been carried out to reach an understanding of how a knowledge net can be designed. In the first study personal home pages on the World Wide Web were explored to see what information people include and how such pages are used. This was a first step in finding what information people would include and/or find useful in a knowledge net. In the second study a prototype of parts of a knowledge net was implemented and evaluated with users. Topics could be entered and related to each other, and one's knowledge about a topic could be rated. One idea behind this study was to investigate if it is possible to describe the knowledge a person has in such a way that a context for the knowledge is included and, thereby, making it possible to find the person most suitable for the task. The third study was an ethnographic investigation of knowledge sharing in an organisation.

Among the findings were that on personal home pages the subjects were willing to present themselves in a way that is of interest for a knowledge net. The subjects tended to be willing to keep their home page up-to-date. Topics entered into the knowledge net prototype were on a rather high level of abstraction. When rating their knowledge about topics, the subjects wanted to be able to specify what parts of the topic their knowledge concerned. Thus, the rating of knowledge seems to help people specify what they know. The software designers that took part in the third study had a rather local contact net that they used frequently. A knowledge net could be one way of extending the number of "experts" to contact.

The results suggest that the knowledge net approach is useful and realizable, and that it in many cases may facilitate the management of knowledge within an organisation. There are still aspects of how a knowledge net should be designed that need to be further investigated, e.g., how one can search the knowledge net and how the results can be presented. How a knowledge net can be supported by management in an organisation is also an important issue.

Sammanfattning

Kunskap anses idag vara den mest värdefulla tillgången i många organisationer. Speciellt stora organisationer kanske vet vilka kompetenser de har men inte var i organisationen eller vem som har dem. Om kunskapen i en organisation kan göras tillgänglig inom organisationen så skulle kunskapen kunna användas mer effektivt. Många försök har gjorts att lagra individers kunskap för att stödja delandet av kunskap.

Denna avhandling fokuserar på den alternativa idén om ett så kallat kunskapsnät. I ett kunskapsnät lagras information om ”vem-vet-vad” snarare än ”vad”. Fokus i ett kunskapsnät ligger på individers kunskap. En graderingsmetod används för att hitta personer med ”rätt” kunskap (inte nödvändigtvis den som vet mest) om ett ämne. Varje ”medlem” väljer vad hon/han vill bidra med till kunskapsnätet vilket medför en vilja att dela kunskapen med andra. Ett kunskapsnät innehåller också faciliteter för att kunna kontakta de föreslagna ”experterna”.

Tre studier har utförts för att få en förståelse för hur ett kunskapsnät ska kunna utformas. I den första studien undersöktes personliga hemsidor på World Wide Web för att se vilken information personer inkluderar och hur sådana sidor används. Detta var ett första steg i att hitta vilken typ av information som är användbar i ett kunskapsnät. I den andra studien implementerades en prototyp av delar av ett kunskapsnät och utvärderades med användare. Ämnen kunde läggas in och relateras till andra ämnen, och kunskapen om ämnena kunde graderas. En tanke med denna studie var att undersöka om det är möjligt att beskriva den kunskap en person har på ett sådant sätt att ett sammanhang kring kunskapen inkluderas och därmed gör det möjligt att hitta den person som är bäst lämpad för uppgiften. Den tredje studien var en etnografisk undersökning hur kunskap delas i en organisation.

Studierna visar att försökspersonerna var villiga att presentera sig själva på personliga hemsidor på ett sätt som är intressant för ett kunskapsnät. Försökspersonerna tenderade att hålla sin hemsida uppdaterad. Ämnen som lades in i kunskapsnätsprototypen var på en rätt hög abstraktionsnivå. När kunskapen om ämnena graderades ville försökspersonerna kunna specificera vilka delar av ämnet de hade kunskap om. Det verkar som om graderingen av kunskap hjälpte försökspersonerna att specificera vad de kan. Programutvecklarna som deltog i den tredje studien hade ett ganska lokalt kontaktnät som de ofta använde. Ett kunskapsnät kan vara ett sätt att utöka antalet ”experter” som kan kontaktas.

Resultaten visar att kunskapsnät är användbara och realiserbara, och att de i många fall kan underlätta hanteringen av kunskap i en organisation. Det finns fortfarande många aspekter av hur ett kunskapsnät ska utformas som behöver undersökas vidare, till exempel hur en sökning kan ske och hur resultatet av sökningen kan presenteras. Hur ett kunskapsnät kan stödjas på chefsnivå i en organisation är också en viktig fråga.

Acknowledgements

There are many persons I would like to thank for their support and contributions to my work. First of all I would like to thank my two supervisors, Professor Kerstin Severinson-Eklundh, who has supervised my research from the start, and Professor John Bowers, who became my co-advisor a couple of years ago. They have supported me with ideas, helped me with my studies, and read and commented on my reports. A third person who has been of help is my husband, Johan. He has read all my reports (often more than once) and supported me with many valuable comments.

I would also like to thank my room mate, Kai-Mikael Jää-Aro, who is always ready to answer my questions about UNIX, LaTeX etc, for valuable comments on this thesis. Ovidiu Şandor and Cristian Bogdan have helped me with Java and JML problems in my work with the prototype of a knowledge net. They have, together with the rest of my colleagues at IPLab, been valuable "parts" of my personal knowledge net!

Last but not least, I would like to give many thanks to all persons who have taken part in my studies. They have all contributed in an important way to my study of knowledge nets.

This work has been funded by the Swedish Transport and Communications Research Board (Kommunikationsforskningsberedningen, KFB) and by the Swedish National Board for Industrial and Technical Development (Nutek).

Contents

1 Introduction **1**
Research questions 2
Overview of the thesis 2

2 The importance of knowledge in an organisation **3**
Knowledge sharing within organisations 5
Knowledge management 7
Conclusion 8

3 Making knowledge available by storing the knowledge **9**
Organisational memory systems 9
Possible limitations associated with storing organisational knowledge . . . 11
Conclusions 12

4 Making knowledge accessible by storing who-knows-what **13**
The knowledge net approach 14
The use of knowledge nets 15
Related work 16
Conclusions 19

5 Study I: Personal home pages on the Web **21**
Introduction 21
Background 21
Method 23
Results 24
Conclusions 28

6 Study II: Letting individuals within an organisation describe their knowledge **31**
Introduction 31
Method 32
Part I: Researchers 33
Part II: Software developers 37
Results 41

Conclusions 48

7 Study III: Studying knowledge sharing between colleagues in an organisation 51
Introduction 51
Related work 51
Method 52
Results 53
Conclusions 61

8 Design aspects of knowledge nets 63
Using personal home pages as, or as the basis for, a knowledge net 63
A higher level of complexity in a knowledge net 64
Using the AETHER Awareness Engine in a knowledge net 68

9 Summary and conclusions 73

A Personal home pages on the Web 81
Interview questions 1996 81
Interview questions 1997 83

B Analysis of the results from study II 85
t-tests 85

Chapter 1

Introduction

While working as a development engineer at Ellemtel (today Ericsson UAB, AXE Research & Development), I became aware of how difficult it was to find information about different matters in a large organisation. Where would I start looking for information? Often the most effective way was to find someone who knew more about the topic and ask that person. If this did not work, the next step was to try to find some documentation that could help. This could be both time consuming and complicated, and might still leave me with unanswered questions.

In the summer of 1994 I started my PhD work which took its starting point in the ideas about "TheKnowledgeNet" (Marmolin, Sundblad, Lidbaum & Tobiasson 1991, Marmolin, Sundblad, Tollmar, Avatare & Eriksson 1992, Marmolin & Sundblad 1993). The main idea behind TheKnowledgeNet was to support people in a knowledge organisation to find other persons with specific knowledge about certain tasks.

The advent of the Internet, and later on the Web, has given not only new possibilities to search for information, but also to distribute information within an organisation or even between organisations. As a result of this, the amount of available information has increased during the last years. Some kind of facility to narrow the search field, either when looking for documented information or when trying to find other persons that can help, would be helpful, i.e., some kind of support that could direct a person to what or whom s/he is looking for.

Within the CSCW (Computer Supported Cooperative Work) area (cf. Bannon & Schmidt 1991, or Dix, Finlay, Abowd & Beale 1998), systems that could support this search for knowledge and information are often based on storing the knowledge or information in a shared information space. This approach can be useful in many situations, but there is reason to believe that the possibility to find other persons, known or unknown, with a certain competence can be even more useful.

My work concentrates on what I call a knowledge net, a general concept for a network of references to persons with certain knowledge. The work relates to the field of CSCW.

Research questions

Since a knowledge net should consist of information about who-knows-what the employees in the organisation, in some way, need to contribute with such information. Much information about what people are working with can probably be found stored in documents within the organisation but this may not be sufficient. The information about who-knows-what has to be presented on a sufficient level of detail and has to be kept up-to-date, otherwise the knowledge net will not be used.

The questions focused on in this thesis are:

- Can a knowledge net be at all useful?
- What type of information should be included in a knowledge net?
- How can information about who-knows-what be entered into a knowledge net and be kept up-to-date?
- How can the "expert" most suitable for a certain question (not always the one that knows most!) be suggested in order to minimise the workload on persons with expertise knowledge?
- How should a knowledge net be designed to minimise the work of entering and changing information and maximise usability?

Overview of the thesis

In Chapter 2, the importance of knowledge in an organisation is discussed. Also, the area of knowledge management is briefly described. In Chapter 3, applications that store knowledge in order to make knowledge available within an organisation are discussed. Chapter 4 discusses the approach to store references to a person with knowledge in order to make knowledge accessible within the organisation, early ideas about knowledge nets, and how knowledge nets can be used.

The studies that have been conducted as part of this thesis are described in Chapters 5, 6, and 7. Study I, in Chapter 5, concerns personal home pages on the Web. People from three organisations were interviewed and their home pages were examined each year during three years. In study II, Chapter 6, a prototype of parts of a knowledge net was used by persons from two organisations. They could enter topics they were interested in or wanted to inform others about, rate their knowledge about the topics, and relate topics to each other. In study III, Chapter 7, a group of software developers were observed and interviewed about how they find persons to ask about different matters.

In Chapter 8 different design aspects for the knowledge net approach are discussed. The last chapter presents a concluding discussion of the knowledge net approach, the research questions, and the results from the studies. Also, future work is described.

In Appendix A the interview questions from study I, about personal home pages, are presented. In Appendix B the analysis of the results from study II is presented.

Chapter 2

The importance of knowledge in an organisation

What is knowledge? This is a rather philosophical question and could involve a discussion of the sociology of knowledge (Coulter 1989) or the difference between intelligence and knowledge (Ryle 1949). However, here the concept of knowledge will be discussed more as a "commodity" to be shared in an organisation.

Davenport & Prusak (1998) distinguish between *data*, *information*, and *knowledge*. Data are only objective facts about events. The data tell nothing about *why* and they cannot predict *how*. The event of going to a store, buying milk, cheese etc, can be stored as data. However, the data cannot reveal *why* that particular store was chosen, *why* those specific things were bought, *how* likely it is that we will go back to the same store, and so on.

Information can, according to Davenport & Prusak, be seen as a message from a sender to a receiver, either in form of a document or as an audible or visible communication. It influences the receiver's judgement and behaviour and, therefore, changes the way the receiver perceives something. A message that contains information for one receiver might only contain data for another, i.e., the receiver, not the sender, decides whether the message is information or not. Information has a meaning and data have not.

Knowledge, on the other hand, is described by Davenport & Prusak as a mixture of various elements:

> *"Knowledge is a fluid mix of framed experience, values, contextual information and expert insight that provides a framework for evaluating and incorporating new experiences and information. . . "*

Knowledge can be derived from information. Knowledge comes from *comparing* one situation with another, identifying *consequences* of decisions and actions, *connecting* one bit of knowledge to others, and from *conversation* with other people about this information.

Gundry & Metes (1996) view knowledge as a personal capability like a skill or experience. It is a capability to do or judge something, today or in the future. They also argue that what is read, seen, heard or felt is not the knowledge, it is the medium used to transfer knowledge, i.e., the knowledge is transferred through the information we receive using a medium. The capability itself can be acquired by using the medium.

Knowledge has, in the area of cognitive science, been divided into two parts; declarative and procedural (Anderson 1990, Wærn & Wærn 1984). The declarative knowledge is about *what*, and has to do with facts. It is knowledge that we are aware of and that we are able to articulate. The procedural knowledge has to do with *how*. It involves procedures of how to create things, i.e., what you do to solve a certain problem, what procedures you perform in order to achieve a result. It is knowledge that guides our actions and how we make decisions, typically outside the scope of consciousness. Some of the procedural knowledge concern things we are unaware of, i.e. riding a bicycle. The procedure of how to ride a bicycle can be described in a set of rules, e.g., if you start falling to the left then turn the handlebars to the left and so on. However, these rules are not enough for a person to learn how to ride a bicycle. There are a number of other aspects that need to be considered when actually performing the procedure and that cannot be formalised as rules (Polanyi 1964). The knowledge contained within these aspects is said to be *tacit* and is in many cases difficult to verbalise.

According to Davenport & Prusak tacit knowledge is almost impossible to reproduce in a document or database. They also argue that in order to store knowledge about an action it is necessary to document rules about how to perform the action. With tacit knowledge it is almost impossible to separate the rules from how an individual acts, and how an individual acts is difficult to describe in words. One example given by Prusak is when he was a child and tried to improve his baseball playing. His father gave him a book on the topic. However, even though the book was written by a skilled baseball player, Prusak still did not significantly improve his playing after reading the book. The skills described in the book were either too complex to be expressed in words or too difficult to learn only by "reading".

Another example of the difficulty to "learn by reading" is Collins' (1985) study of several attempts to replicate the TEA-laser (Transversely Excited Atmospheric laser) in the beginning of the 1970s. He found that "no scientist succeeded in building a laser by using only information found in published or other written material". All successful replications of the laser were due to knowledge acquired from personal contact and discussion. He also found that "no scientist succeeded in building a TEA-laser where their informant was a 'middle man' who had not built a device himself", and also that an extended period of contact with the informant was necessary for success. This illustrates the importance of personal contact and communication even when printed documents are readily available.

Degerblad (1988) differentiates between explicit and tacit knowledge in that explicit knowledge can be shared through the use of language, either written or spoken, and that tacit knowledge can only be shared by taking part in an activity. He also argues that there is a technical and a normative side of knowledge. The technical side of knowledge is theoretical, and involves the "knowledge to make statements" and the knowledge to use facilities available in the area. The normative side of knowledge is the kind of

knowledge that involves good and bad examples about the area, i.e., knowledge that reveals that we know the area knowledge that makes us able to make judgements within the area. Explicit knowledge can be, on the technical side, to be able to propose and, on the normative side, to know the area. Tacit knowledge can be, on the technical side, to be able to use facilities in the area and, on the normative side, to be able to make judgements.

Related to the term knowledge is the term expertise. Expertise, knowledge that has been developed over time (Davenport & Prusak 1998), demands a very careful consideration before being stored and reused by other persons. Otherwise, it is easy to loose important parts of the knowledge or even to misinterpret the knowledge. Ford, Bradshaw, Adams-Webber & Agnew (1993) argue that "expertise is not a natural resource that can be harvested, transferred, or captured".

Knowledge sharing within organisations

There are, according to Davenport & Prusak, three factors that contribute to the inefficiency in how knowledge is managed in organisations. The first is that the information about where to find knowledge is often incomplete. This makes it difficult to find the knowledge, if it can be found at all. The second is that the same knowledge resides in many different places and at different levels of detail. This makes it difficult to know which source is the most suitable one. The third is that people rather ask a person in the office next door than try finding someone elsewhere in the organisation that may be better suited to answer the question. To be content with too low a level of knowledge may lead to an unnecessary sub-optimisation when solving a problem.

The task of finding persons with a certain knowledge is supported by a knowledge net. The main idea with a knowledge net is to make it easier to exchange knowledge in an organisation and to make it possible to find persons with knowledge relevant in a specific situation. Why then would it be better to find a person? Why not just store all knowledge in a database and everyone can find whatever information they are looking for by searching the database?

Bannon (1986) argues that people rather ask other persons for advice than search through a manual for information. He found, when interviewing administrative and clerical personnel, that the major source of information about the computer systems used were other users. One person in his study expressed that sharing office with a person more experienced within a certain area provides an ideal environment for solving problems related to the area. Bannon also points out that a common view of the problem facilitates an understanding and solving of the problem, especially in the case of a novice-expert conversation (see also Clark 1996, Chapter 1). He argues that in a face-to-face conversation interruptions can provide feedback about the participants' understanding of the current dialogue, and the conversation can change to an appropriate level of understanding. Thus, to meet a person face-to-face could help solving a problem.

There is also a difference between formal and informal sources of information. A formal source can, e.g., be a computer system help desk while an informal source can

be a person not officially with the task of helping other persons. The reason that formal sources often fail in their mission is because the persons working at a formal source do not know enough about the particular topic and because they often are remotely located (Bannon 1986). Instead, informal sources such as colleagues are chosen because they are physically available, they are personal friends, or they are known to be experts on the topic. Investigations show that people working in software design projects prefer to ask nearby colleagues rather than using formal information sources (Waterson, Clegg & Axtell 1997, Eveland, Blanchard, Brown & Mattocks 1994). The reason is that the colleagues better relate the question to the problem.

Another approach on how to use formal sources is when domain experts are integrated in a development group. This approach was studied by Gantt & Nardi (1992) in the use of CAD systems. They found that, in some of the workplaces, local developers, i.e., domain experts with advanced knowledge of computing, had been given the official role to help the end users of the CAD system. This shows that even though formal documents is available it still is regarded as useful to provide local experts.

In software development projects it is, according to Kraut & Streeter (1995), often the case that many persons working on the project do not have sufficient knowledge about the domain they are working in, and that information needed to make decisions in the project is not available through documents. Kraut & Streeter found that other people were the most used and valued sources of help in software development projects. The value given a person's knowledge was higher than what could be inferred from the actual use of that knowledge. Also, the ease of getting the information is a critical determinant for asking other persons. Even though some people were difficult to access, e.g., people outside the group or even the company, and used as information sources more seldom than other people, they were still judged to be valuable sources of help. All forms of written documentation were valued less than personal contacts.

In a study of collaboration within a research environment Kraut, Egido & Galegher (1988) found that researchers who have offices next door to each other communicate approximately twice as much as those whose offices are on the same floor, but located more distantly. They also argue that although one explanation is that people with interests in common often are located closely to each other, people are more likely to get acquainted and identify shared interests if they get the opportunity to meet. This was likely to occur around the lunch table, in a corridor etc (Kraut, Galegher & Egido 1987–1988). This illustrates that the location is important when communicating with other persons. It is not only important to be located in offices nearby but also to get the opportunity to occasionally meet.

Not only colleagues who are part of the same group, participating in the same project or sitting nearby, can help solving problems. Also people outside the group are found to be important when searching for information (Kraut & Streeter 1995).

The management's attitude plays a central role for the sharing of information among persons in the organisation (Bannon 1986, Davenport & Prusak 1998). In a field study of an organisation within the U.K.'s central government using a local area network with CSCW-related applications (Bowers 1994) it was found that the members of the organisation did not use the monthly report application in an appropriate way. Therefore, the managers changed the reporting procedure to having each member of the organisation

report activities when they occurred, instead of once a month. This was perceived as a lack of trust in how the members spent their time, and the members, therefore, refused to use the procedure. People must feel that they can trust the organisation and that they will be rewarded when sharing their knowledge with others.

Knowledge management

Knowledge Management (KM) can be defined as "the process of identifying, capturing and leveraging knowledge within an organisation and using that information to increase profitability and competitive advantage within the marketplace" (Densford (1996) quoted in Hollocks (1998)).

Knowledge management concerns knowledge that resides in an organisation; how it can be identified, captured and leveraged to be used today and in the future. Kühn & Abecker (1997) have identified several important aspects that concern knowledge management. Such aspects are the time spent on looking for needed information, the number of valuable pieces of information that is buried in piles of documents and data, the number of employees with essential undocumented knowledge, the repetition of costly errors due to lack of knowledge of previous experience, and the insufficient information flow which results in product delays and suboptimizations. These aspects are good examples of areas of interest for knowledge management.

Benjamins, Fensel & Pérez (1998) describe four actions that are involved in managing knowledge. The first is knowledge gathering—to collect the knowledge that is to be managed. The second is knowledge organisation and structuring—the knowledge collected is given a structure in order for effectively managing. The third is knowledge refinement—maintaining the knowledge to keep it up-to-date and correct. The fourth, and last, is knowledge distribution—to bring the knowledge to whoever needs it. All these steps need to be considered when computer tools to support knowledge management are developed.

A broader description of the knowledge management area is given by Stewart (1997*b*). He describes three important assets of knowledge in an organisation. One is the knowledge that each person working in the organisation has, called the human capital. The second is the knowledge of individuals that is stored and used by others in the organisation, called the structural capital. The last source of knowledge is the relations to the customers that are interested in buying whatever the organisation is selling, called the customer capital.

The structural capital can, to some extent, be stored in some kind of shared information space to be available to the whole organisation. A shared information space that stores individuals' knowledge is, within the area of CSCW, often referred to as an organisational or corporate memory.

The distinction between information and knowledge is, according to Gundry & Metes (1996), important in knowledge management tools. They point out that such tools do not actually manage knowledge. Instead, they help capture, organise, store and transmit information from which an individual may acquire knowledge.

Benjamins et al. argue that there are two types of knowledge management systems. There are vertical systems that are developed inside a company and for a specific purpose. Such systems have little value for other situations than for the ones it was developed for. There are also horizontal systems that are more general and can be applied to a variety of situations. The knowledge net approach can, according to this "classification", be considered a horizontal knowledge management "system". On the other hand, an application of a knowledge net may, or may not, be a vertical knowledge management system.

Conclusion

There are different ways of looking at knowledge. As can be concluded from the discussion in this chapter knowledge is difficult to capture, it is complex, and it is dependent on the context in which it was captured.

The need to interact with some experienced person is, in many cases, necessary in order to succeed with a task. To ask persons with a shared common ground of the area is important in order to avoid misunderstandings. Efficient systems and methods to support the interaction with other, experienced, persons are important. This will most certainly save time and, thereby, money within organisations.

To enhance the knowledge sharing in an organisation the "organisational" information or knowledge, i.e., information about the organisation and knowledge that individuals in the organisation have, can be stored and/or the information about who-knows-what can be stored. Applications that store organisational information are often referred to as organisational or corporate memory systems. Such applications will be discussed in more detail in the next chapter. Applications that store references to who-knows-what will be discussed more in Chapter 4 together with the knowledge net approach.

Chapter 3

Making knowledge available by storing the knowledge

The interest in CSCW applications (or groupware) has steadily increased during the last decade. Problems when developing such applications have been many but, nevertheless, a number of CSCW applications are used with success today, e.g., email and calendar systems (Grudin & Palen 1995, Bowers 1994). Grudin and Palen argue that the reason why CSCW applications are more successful today than ten years ago, when Grudin (1988) performed a similar study, is that today organisations have common hardware platforms that support such systems, that they provide strong technical support to install and maintain such software, that people are more comfortable with computer technology today, and, finally, that computer products of today are more powerful.

Lately, it has become common to use the Web as a platform for CSCW applications (Trevor, Koch & Woetzel 1997, Dix 1996). One example is the BSCW (Basic Support for Cooperative Work) system (Bentley, Appelt, Busbach, Hinrichs, Kerr, Sikkel, Trevor & Woetzel 1997). BSCW provides a number of facilities for creating shared workspaces, i.e., repositories where shared information can be stored in an organisation. Information such as documents, member contact information, pictures, URL links to Web pages, threaded discussions and so on can be stored in the shared workspace. A shared workspace can be accessible to the members of a group using a user name and password. The BSCW system manages workspaces for different groups, and users can be members of more than one group.

Organisational memory systems

Many collaborative systems that support sharing of knowledge within an organisation store the knowledge in a shared information space, they are often referred to as organisational memory systems. The information can then be retrieved and used in the future.

Terveen, Selfridge & Long (1993) developed a tool that "serves as a 'living design memory' for a large software development organisation". The purpose of the tool was to

store and, if necessary, evolve design expertise when developing software applications. The tool consists of a "design knowledge base" that records relevant information, and a "designer assistant" that provides access to the knowledge base.

Design knowledge that was identified as important during the design process but that was difficult to find was, e.g., properties of the current implementation, real-time and performance constraints, and the impact of design decisions on other aspects of the software. The authors found that those developers who had effective expertise networks and thus knew who to talk to about these questions were those that were most successful. The authors argue that to store this kind of information and to make it change over time ("living", which emphasise that the system "must evolve in response to problems detected with it or changes in the knowledge situation of the organisation") would give all developers the opportunity to find information relevant for their work.

In another organisational memory system, Answer Garden (Ackerman & Malone 1990, Ackerman 1994), commonly asked questions about an application domain are stored, together with the answer, in a common repository. The user browses the tree structure for a certain question. If the question is not found then there is a possibility to email the question to an anonymous expert who then answers the question and adds both the question and the answer to the repository. This kind of organisational memory system may help people in their search for facts about specific application domains.

In a new version of Answer Garden, called Answer Garden 2 (Ackerman & McDonald 1996), the facility to find help has been extended, and there is also a possibility to customise software components in the system. In Answer Garden 2, the experts answering questions can be either anonymous or known by the user. A question can also be directed locally, to someone nearby. If this does not yield an answer then the question can, by an "escalation agent", be directed to a global expert. This version of Answer Garden has been extended in the part where people can direct their question to an expert, and is, in some regards, a knowledge net about one specific domain.

Yet another organisational memory system is the gIBIS system (Conklin & Begeman 1988). gIBIS is a graphical implementation of the IBIS, Issue Based Information Systems, method (Rittel & Kunz 1970). The IBIS method focuses on key "issues" in design problems. Each issue can have many "positions", which can be a statement or assertion which resolves the issue. A position, in turn, can have one or more "arguments" which either support or object to that position. The issues, positions, and arguments are presented in a hypertext tree structure. The purpose of the gIBIS system was to explore the capture of design history, to support computer mediated teamwork, and to investigate hypertext navigation of very large information spaces.

QuestMap (Conklin 1996) is a "computer tool for capturing and managing any size of IBIS map, and any number of interlinked maps, among large number of users". QuestMap is used to capture the key points in meeting conversations and to show these during the same or other meetings.

All these different kinds of organisational memory systems also reflect the diverse use of the concept organisational memory. Walsh & Ungson (1991) describe the concept as "it refers to stored information from an organisation's history that can be brought to bear on present discussions". They argue that the core of an organisational memory is formed by information about decisions made and problems solved over time. They also

describe five possible locations where acquisition and retention of organisational memory can take place. The first is *individuals* who "retain information based on their own direct experiences and observations". The second location is *culture* which "embodies past experience that can be useful for dealing with the future". The third location is *transformations* which embody information about "the logic that guides the transformation of an input (whether it is a raw material, a new recruit, or an insurance claim) into an output (be it a finished product, a company veteran, or an insurance payment)". The fourth location is *structures* which concerns "individual role behaviour and its link with the environment", i.e., the individual role provides a repository where organisational information can be stored, and the role concept provides a link between individual and organisational memories. The last location is the *ecology* of the workplace of an organisation that "encodes and thus reveals a good deal of information about the organisation".

Bannon & Kuutti (1996) point out that a definition of organisational memory as above, with an attempt to include almost everything, leaves one wondering what, within an organisation, is not part of an organisational memory.

The articles discuss the variety of definitions of organisational memory, a concept used in many different disciplines. It appears that there are still lacking a consistent and definite definition of what we mean by organisational memory.

Possible limitations associated with storing organisational knowledge

Although systems that store knowledge may be useful in many situations there are also many unsolved problems, such as:

- how can the knowledge, together with the context, be formalised and expressed in a form suitable for storing,
- how can, if necessary, the stored knowledge be modified over time to reflect changes in the environment in which it was stored and is to be used,
- how can a changing, growing and evolving database be designed to be accessible and scalable?

As pointed out by Bannon & Kuutti (1996) and Bannon & Bødker (1997) it is difficult to predict what knowledge, or information, within an organisation, will be of interest in the future and thereby is worth storing. This also involves a tradeoff between the cost of storing and the cost of reinventing, but how is this tradeoff to be decided? Bannon & Kuutti also argue that if the "activity" during which the "storing" takes place differs from the one in which the "remembering" takes place, then the information may be re- or even misinterpreted. It is a question of what needs to be stored and what can be left as taken for granted.

Bowers, Viller and Pycock (in preparation) have observed users of an organisational memory system "creatively" classify events and knowledge in order to anticipate future uses and target future potential users. What the users did was to work out rules and

recommendations of collected experience, and store these in various databases. How the users did this was affected by their own knowledge about the facts. To store knowledge together with an activity is an interesting approach (see further Chapter 6 on rating individuals' knowledge).

Ellis, Gibbs & Rein (1991) discuss the importance of organisational knowledge when developing groupware. Organisational knowledge is, according to Ellis et al., knowledge about an organisation's structure, history and goals. They point out that this knowledge is volatile and difficult to capture. They also mention group knowledge, which, as well as organisational knowledge, is important when specialising a groupware tool to meet a group's particular needs. Group knowledge could for example be user and group profiles describing particular needs of the group.

Not surprisingly, it seems like storing knowledge does not solve all problems concerning sharing knowledge in an organisation. People also need other persons to discuss problems with. An example is "garment technologists" (garment technologists specify particular aspects of garments that suppliers are to manufacture) working in the fashion industry (Pycock & Bowers 1996). To create a specification of garments the garment technologists often used an old specification in discussion with other colleagues "as a resource for suggesting past examples", i.e., they used other persons as a source of knowledge. Pycock & Bowers also found that many "new" problems had actually already been solved earlier. It appears that, in an organisation, there exists knowledge that may be unknown but useful.

Conclusions

It appears that it is difficult to store knowledge without loosing or "changing" the context, and thereby, loosing an important part of the knowledge. The context is necessary in order to ensure that the "knowledge" stored retains its original meaning.

Chapter 4

Making knowledge accessible by storing who-knows-what

An approach other than storing the knowledge itself in a database, as in organisational memory systems, is to store references to persons with the knowledge, i.e., references to "who-knows-what", instead of "what", is stored.

When using a system that stores references to who-knows-what there is no need to decide what information will be relevant in the future. As long as the person with the knowledge remains within the organisation the knowledge will be available and also (hopefully) "automatically" updated. Also, the knowledge and the context around the knowledge do not have to be described before being stored in a database, since the knowledge when sought for will be mediated by a person (the "expert"), i.e., in such a system the context is not separated from the knowledge. Neither is the problem with formalising "tacit" knowledge an issue in such a system, because that kind of knowledge need not be stored at all. Finally, the amount of data that need to be stored is much less in a system that stores references to who-knows-what instead of storing the knowledge itself.

Bowker (1996) argues that people cannot remember the past situation correctly, and that the present is used as a benchmark. He also points out that "total recall, in individuals or organisations, is neither desirable not possible". Bowker argues that there are several good reasons to forget things about the past, e.g., that rediscovery is easier than remembering, and that forgetting facilitates the change of identity of an organisation. He further argues that the use of memory as a tool for reification can have harmful consequences.

The idea to store references to who-knows-what has been described by, e.g., Stewart (1997*a*), Stewart (1997*b*), and Davenport & Prusak (1998). Stewart argues that an extended "telephone book" containing not only name and contact information but also information about what competencies the person has could be useful in a knowledge organisation. He calls this "in-house, or corporate, Yellow Pages". Davenport & Prusak

use another metaphor. A “knowledge map” is a “map” of what knowledge resources, e.g., persons or documents, that exist in the organisation and where they are located.

The knowledge net approach

“In-house yellow pages” and “knowledge maps” can be seen as two names for the general concept of storing references to knowledge instead of knowledge itself. The generality of this concept makes it rather difficult to create general design rules for an application based on the concept.

In order to be able to make a general design proposal of an application where the ideas of storing references to knowledge are used we use a more “restrictive” and, as we believe, more usable description of such a system. We call this a knowledge net (Groth 1997). In the knowledge net approach we focus on people. A person with a question or a problem should be able to get an answer or help from another person. We believe that to interact with another person (instead of being referred to documentation) gives many advantages.

Another main principle of a knowledge net is that the knowledge a person has is given a rating. The rating is a description of what and how much a person knows about a certain topic. What a person knows about a topic makes it possible to detect the person with knowledge about the relevant aspect of the topic. How much a person knows about a topic makes it possible to detect the right level of competence. This is necessary in order to find the “right” person, which not necessarily is the one that knows most. This is important in order to minimise the work load on the “experts”.

A third principle is to let each individual describe their knowledge, either automatically or by hand, and be responsible for it (compare the information on personal home pages). This is important because we believe that if people can choose what information to supply then this is also information that they are willing to share with others.

A fourth and last principle is that there must be some way of contacting the person(s) being referenced either in person, or by email, telephone etc. The medium used depends on who is referenced and what preferences that person has.

When using a knowledge net the user is directed to a person with knowledge about the topic. The user must also be given advice on how to contact this “expert”. The awareness of how the person can and wants to be contacted in general and at this moment is important and some kind of tool to enhance the communication between persons is an advantage. An ideal solution would be to have an awareness application (cf. Dourish & Belotti 1992, Dourish 1993, or Tollmar & Sundblad 1996) integrated with the knowledge net application. Such an application could give information about the availability of the “expert” and how the “expert” can be contacted.

The first ideas about a knowledge net-like system were presented as a part (which, however, never was completed) of a prototype of a computer supported cooperative work environment called CoDesk (Marmolin et al. 1991, Marmolin 1991, Marmolin et al. 1992, Marmolin & Sundblad 1993, Tollmar, Marmolin & Sundblad 1994, Tollmar & Sundblad 1995). The vision of “TheKnowledgeNet” within CoDesk was “a distributed ‘library’ of documented and undocumented individual knowledge that is made

available to all team members by communication" (Marmolin 1991). TheKnowledgeNet was to be integrated with CoDesk. Together they would "support a social work situation in which collaboration among peers can take place by sharing and integration of knowledge" (Tollmar & Sundblad 1995). One primary idea with TheKnowledgeNet was to let people within the collaboration environment describe their own knowledge areas.

The work with TheKnowledgeNet was much about "what" and not so much about "how". There were many ideas of what parts TheKnowledgeNet should consist of—a set of requirements on the application. How these requirements would be implemented and used was not dealt with. Therefore, my focus on the work with the knowledge net approach has been more on how facilities in the application of a knowledge net can be implemented and used. Also, since the CoDesk environment is no longer an issue, the requirements of an application such as TheKnowledgeNet has changed.

Difficulties with applications that store references to who-knows-what appear if people disappear from the organisation. This weakens the system. Another difficulty is that it does not suffice to store a reference like "John knows about Java", because it does not reveal *why* John is a suitable expert for a question about Java. In a large organisation there could be several people that know about Java, but who is most suitable for a certain type of question? Therefore, it is also necessary to have some kind of description of how much John knows about Java and also what parts of Java he is familiar with. This is probably one of the most challenging parts of such a system because the description of what it is the person knows about the topic, together with information about the organisation, should make it possible to detect the person most suitable to ask about a certain matter. There is also a balance of not storing too much and storing enough. If too much is stored then the problems identified with organisational memory systems will also be a problem here. On the other hand, if too little is stored it can be difficult, or impossible, to identify the most suitable person to ask.

The use of knowledge nets

The main idea with a knowledge net is to support persons in finding who-knows-what, i.e., other persons with a specific competence, in the organisation. Possible users of a knowledge net within an organisation are

- persons who have recently been employed, and who need to develop a personal contact net of whom to ask about different matters,
- persons who seek a certain competence to discuss a problem with but do not know where to find it in the organisation,
- "isolated" persons and persons who are part of a virtual organisation or an informal network, who seek additional competencies,
- project managers who need to find the right competence for a new project,

- higher level managers who need to see what competencies are missing in the organisation,
- other persons who, e.g., by telephone, direct people to persons within the organisation such as receptionists, or who help customers solve different problems such as help desk staff.

For some purposes a knowledge net can be built and supported by employees dedicated to this task. This can be the case in, for example, consultant firms where each consultant's profile is important when selling competence to customers. A profile of each consultant is needed. The profile could describe, e.g., which projects and specific tasks the consultant has been involved in. These profiles could then be supported by a specific department and used by, e.g., managers who want to find a person with the right competence to work in a project.

For other purposes it could be in each person's interest to build and support references about oneself. This can be the case when, for example, several individual consultants share a business. In this case each individual consultant is most likely responsible for his or her profile. The profile could be used, e.g., by a receptionist to direct a customer to the consultant with a competence suitable for the task.

Something in between these two examples can be larger knowledge organisations. Here each individual's competence can be important in several parts of the company, but there are too many employees to make it possible for any person to know exactly who-knows-what. In this case each person's profile could both be generated in an automatic way, e.g., through some information retrieval process, and manually by the person him- or herself. It could be used by other employees within the organisation to find the right person to ask or discuss a problem with.

Related work

Work related to the knowledge net approach has been described in the literature. One is the Microsoft SPUD project, described by Davenport & Prusak (1998), another is the REFERRAL WEB developed by Kautz, Selman & Shah (1997), and a third is a simple kind of knowledge based system called a reference system described by Kjellin (1994).

Other examples of work, not further described here, that is related to the knowledge net approach are using ontologies to create relations and finding persons with knowledge (Benjamins et al. 1998), and the Knowledge Garden which is a Virtual Reality (VR) world where people can find other persons to ask (Crossley, Davies, McGrath & Rejman-Greene 1998).

The Microsoft SPUD project

Davenport & Prusak describe a project at Microsoft, Skills Planning "und" Development (SPUD), where a "people oriented knowledge map" has been developed.

In the SPUD project four types of knowledge structures are used to evaluate people's competencies. These are foundation knowledge, local or unique knowledge, global

knowledge, and universal knowledge. Within each of these four knowledge structures there are two different categories; explicit and implicit competencies. Explicit competencies involve expertise in specific tools or methods while implicit competencies involve more abstract thinking and reasoning skills. Within each type of knowledge competency there are also four defined skill levels; basic, working, leadership, and expert. Exactly how these skill levels are defined is not revealed but they are claimed to be clear and easy to measure. However, the skill levels are interesting because they seem to tell how a person is using her/his skill. The "basic" level is used when a person has started to learn about the area. The next step is when the person uses the skill in a real work situation, "working". A next step would be that the person has responsibility for the area where such skill is used, "leadership". An "expert" level is usually something that requires an acknowledgement from the top level management of the organisation.

Each employee is evaluated in terms of what knowledge they have exhibited in the projects they are working with. Each employee initially rates, together with his or her supervisor, the knowledge, but eventually the entire work team participates in the rating procedure.

This knowledge map is used in the whole organisation. An example of a query given by a user of the knowledge map is "Give me the top five candidates who have leadership skill levels on 80% of the knowledge competencies for this job and who are based in Redmond". It would be very interesting to see how the categorisation works from a user perspective and how well the given references meet the requirements of the user.

Searching documents written by experts

Kraut & Streeter (1990) describe a system that helps identify relevant experts. This is accomplished by comparing an information query with different documents written by experts. The system returns a list of experts, ranked in order of how many relevant documents they have recently written. How this ranking reflects a correct level of competence is difficult to understand. The authors claim that such a system works because no extra effort is needed from the experts in order to enter information into the system—the documents used when comparing a query are written as part of the experts' normal work effort. Also, using these documents makes the system cover the whole company, the user does not need to read all these documents, and, finally, the matching algorithms used when comparing a query with documents do not require that the expert and the person writing the query use the same terminology.

This is an interesting application that uses an advanced information retrieval method on documents when finding "experts". The method could be of interest to use in a knowledge net when searching documents for information to be entered. However, the application lacks features that can detect the person most suitable to ask.

The REFERRAL WEB

The REFERRAL WEB is an interactive tool on the Web that helps people find short referral chains between themselves and experts within a certain area. It uses publicly

available Web pages to create a referral chain. No information needs to be entered by the users.

A referral chain is created by searching for names and following links on Web pages. If two or more names occur in close proximity on a Web page then this is seen as evidence of a direct relationship between these persons. It can be links found on home pages, lists of co-authors in technical papers and citations of papers, and organisation charts. The REFERRAL WEB uses a search engine to retrieve Web pages.

A user of the REFERRAL WEB can either search for other persons by name and then get a chain between himself or herself and a named individual, or search for an expert by specifying a topic. The user is presented with a graph of names of persons and relationships between these persons, a chain of names. The chain can be used to decide which expert should be contacted, either by asking persons in between or drawing conclusions about peoples' backgrounds.

The REFERRAL WEB is an interesting application for finding persons with specific knowledge. It is used today within a limited community of AI researchers. A demonstration of the application is available on http://www.research.att.com/~kautz/referralweb/ . One advantage with the application is that no one needs to enter any information. Two disadvantages could be that the application does not distinguish between two persons with the same name, and no support is given to help the user detect which person is most suitable to ask.

A reference system of experts

Another type of application similar to the REFERRAL WEB is a simple type of knowledge based system called a reference system (Kjellin 1994). In a reference system a number of objects and relations between objects are collected. This kind of reference system has been used, e.g., at an unemployment office where different companies could add information about job opportunities into the reference system. People looking for a job could use the reference system to find a job suitable for their competence.

Each object in the reference system can correspond to a topic, report, project, course, person etc. Objects can be of different types; general, specific or result. A result object is what is presented as a result of a query.

Each relation is dependent on the type of objects it represents. If, for example, two topic objects are related then the relation can be "isA" or "hasA". Or, if one object is a topic and the other is a person then the relation can be "has_knowledge" or "is_known_by".

Each relation can also have a weight determining how strong the relation between the two objects is. For example, a higher weight on a relation implies that the object at the end of that relation probably is more suitable than the one with a lower weight. Using some kind of "relevance feedback" these weights can be updated automatically depending on how well the reference given was.

When using a reference system like this the user is given a question to fill in. The question is, in this case, the name of the topic that the user wants to find an expert in. When the question is entered (i.e., a topic name is given) the reference system tries to match the given topic with all relations to and from that topic object. The objects of

type person that are found will be given as an answer in an order given by the weight on each relation or path of relations.

In order to make a reference system of experts (i.e., a knowledge net) work in an organisation much effort is needed both from the system and from the organisation. The reference system must be easy to use and easy to support—the effort of using and supporting the reference system must be less than the gain (Grudin 1988, Grudin & Palen 1995). The people in the organisation must be willing to supply information and share it with others in the organisation.

If techniques such as information retrieval and relevance feedback can be used then there are many possibilities to make it easier to use and support a reference system. Information retrieval is used to collect information from written material. The advanced method for retrieving information from documents described by Streeter & Lochbaum (1988) and Lochbaum & Streeter (1989) would be interesting to use. Instead of only matching a query with keywords in documents they create a semantic structure space of all keywords. When matching a keyword the method also finds related keywords which will be included in the answer. Relevance feedback is used to automatically update how well the reference given was. However, these techniques are probably not sufficient when creating a reference system including information about who-knows-what because they cannot distinguish between important and unimportant information and they cannot capture everything, which, on the other hand, is not desirable.

There is also a question of methods such as information retrieval and relevance feedback should be used in a knowledge net application. There is a balance between how much such techniques help and how much they complicate the use of the application. This is further discussed in Chapter 8.

Conclusions

Various approaches can be used when designing systems for who-knows-what. We suggest an approach called a knowledge net, which is characterised by a focus on persons (not documents etc), a rating mechanism for the knowledge, a personal responsibility for the information entered, and, last, a function which allows interaction with the expert. We believe that this definition provides a starting point for developing general design criteria.

To find an appropriate way of designing and using knowledge nets three studies have been performed. In the first study (Groth 1998), described in the next chapter, personal home pages on the Web were studied. In the second study, described in Chapter 6, a prototype of a simple knowledge net was implemented and tested. In the third study, described in Chapter 7, a group of software developers was studied with regard to how they find other persons to ask about certain matters.

Chapter 5

Study I: Personal home pages on the Web

Introduction

What kind of information would people find useful in a knowledge net? Are people willing to share such information about themselves with others? Is it reasonable to believe that such information can be kept up-to-date? These are three fundamental questions in the design of a knowledge net.

On personal home pages on the Internet people usually describe who they are and what they do. Since personal home pages contain information about persons, and are created, in most cases, by the persons themselves they, at least, fulfil two of the four principles characterising a knowledge net. This makes it interesting to see how people view their own and other persons' home pages and if and how they use other persons' home pages. A systematic study of personal home pages could give important input about the design of knowledge nets.

Background

The Internet gives a possibility to move information between computers in different networks. Services on the Internet are, e.g., the Web, email, video conferences, specialised databases and archives with information. The Internet can be used for communication and collaboration in both synchronous (at the same time) and asynchronous (at different time) mode.

Intranets, i.e., "local Internets" reachable only by persons within a group or an organisation, are becoming more and more popular. The information on a remote server located anywhere in the world can be reached from any workstation connected to the intranet.

The Internet-technology is relatively cheap and easy to learn and has therefore become a powerful tool for communication and information sharing among small organisations, groups and even individuals. The Internet has in a sense become the first global information medium that is easily available for everyone. The effects of this have only just begun to emerge.

The Web is a global distributed multimedia service that has rapidly become a commonly used medium for distributing and exchanging information. The Web was invented in late 1990 by Berners-Lee (Berners-Lee 1996*a*, Berners-Lee 1996*b*), as an "internet-based initiative for global information sharing", a shared information space with public and private information. The idea of the Web was actually to "be a pool of human knowledge" and to "allow collaborators in remote sites to share their ideas and all aspects of a common project" (Berners-Lee, Cailliau, Luotonen, Frystyk Nielsen & Secret 1994).

One advantage with the Web is that it is not intrusive. Erickson (1996) points out that an important difference between finding personal information using the Web and using email or phone is that the person searching for information is not obligated to the person giving the information. According to Erickson a kind of social debt is created when one asks for information from somebody—if you ask a person for a paper then that person might expect you to read and comment on it.

A home page on the Web is the intended entry point of a logical information structure (usually called a web site) from which all other pages on the site may be reached, directly or via other pages, by hypertext links. If a home page is written (partly or wholly) by and about an individual, it is called a personal home page.

Personal home pages give an individual user of the Internet, or of an intranet, an opportunity to present personal information to other users. These pages on the Web can, therefore, be seen as a source of knowledge about individuals within a network. The network is either global, i.e., available to everyone on the Internet, or local, i.e., available only to the specific users of a certain intranet (physically these individuals may be distributed).

Why study personal home pages?

According to Tauscher & Greenberg (1997) only a small number of pages on the Web are regularly visited. Their study of 23 Web users showed that 60% of the pages visited were only visited once and only 4% were visited four times or more. However, one category of pages that was frequently accessed was personal home pages. Personal home pages were found to act as a gateway to pages for the organisation or individual. This shows that information on personal home pages are of interest for people.

The use of personal home pages was discussed in a workshop on "HCI and the Web" (Instone 1996). One finding was that personal home pages give an opportunity to find out more about certain topics because people can visit a personal home page of someone who is also interested in that topic. This can be compared to Erickson's (1996) argument that the Web can be characterised as a social hypertext, where the nodes represent people. Personal home pages contain information that reflects the personality and interests of the author of the page, and together become a network of personal information.

Instead of using search engines to find information people can go to a home page of a person they know have similar interests and find links to other interesting places, e.g., organisations or persons.

The focus of the present study was to explore how personal home pages in an organisation were used and if this could have any implications for a knowledge net. In order to investigate this it was of interest to see

- how and how often people use the Web,
- when and why people make their home pages,
- how information is presented on personal home pages,
- what people choose to include on their home pages and why,
- what types of information people find relevant on other persons' home pages,
- to what extent personal home pages create new contacts.

Related work

In another study of personal home pages (Bly, Cook, Bickmore, Churchill & Sullivan 1998) 28 personal home pages were examined in four organisations, and 20 of the authors were interviewed. One of the organisations was a college and those persons' home pages could be reached from the Internet, while the other three organisations used an intranet structure supporting personal home pages.

It was found that 75% of the examined personal home pages contained project related information, 50% contained personal information such as hobbies, and 39% contained photos of the author.

Although most pages were used for project information authors from all four organisations were found to include personal information. This was because the authors wanted to give people an idea of who they are. However, some authors did not want to include any personal information, because they felt it did not belong on the page or because they did not know what to write.

Most of the authors expressed that they value personal information on other persons' web pages, because that kind of information gives an idea of who the person is.

Method

This study was performed using semi-structured face-to-face interviews in combination with a demonstration of the subjects' personal home pages. The interview questions are presented in Appendix A. Notes were taken during the interviews. All in all 22 interviews were performed among three groups of people from different kinds of organisations from late spring to late summer 1996.

The first group consisted of eight researchers (one female, seven male, average age of 35 years) at a university (the "uni" group). The second group consisted of six researchers (three female, three male, average age of 32 years) from a research institute

(the “inst” group). The third group consisted of eight software developers (three female, five male, average age of 32.5 years) at a larger development company part of a world wide organisation (the “dev” group). This company had an intranet accessible from company offices in a number of countries.

The persons interviewed were chosen because they had a personal home page on the Web. No one that was asked for an interview declined to take part. Within the groups all persons, except one person in the inst group, belonged to the same “department” within their organisation.

After one year, in the spring and summer of 1997, seven of the researchers and three of the software developers were interviewed once again about the changes made to their home page. The home pages of all 14 researchers and of three of the software developers (the other five had left the organisation or were absent for other reasons) were also examined.

After yet another year, during the summer of 1998, the home pages of the 14 researchers and one of the software developers (the only person still in the organisation) were once again examined.

Results

Most persons interviewed (19) had used the Web for more than a year. As many as 13 subjects had used it from the beginning of the Web, in 1993/1994. Most subjects (18) used the Web every day. The most common activity was to search for information, but eight persons also used the Web as a platform when developing software.

About half of the persons interviewed (12) had had their home page for more than a year. The reasons reported by the subjects for having made a personal home page were because they found other persons’ home pages useful, they wanted to distribute their publications, they wanted to try out the new medium, and because they considered it a good way to find information about other persons. Four of the researchers pointed out the usefulness of other persons’ home pages. Many (16) of the persons interviewed expressed that they still wanted to make changes to their home page. They wanted to make changes in the content, add links and articles, change the layout or the colours used on the page etc. Nine of the subjects were satisfied with the home page they had. Apparently, people find the information presented on other persons’ home pages to be of interest which indicates that the kind of information presented on such pages could be of interest in a knowledge net.

There were some differences between the groups. One difference was that the dev group had their home pages on an intranet. This could affect the choice of information to include on the home page. Another difference was that the researchers had been told by higher management to have a personal home page. In the inst group pictures had been taken and a home page with the picture and name of the person was created by the organisation. This was not the case in the other two groups.

What kind of information do people include on their home page?

The home pages of the subjects were quite homogeneous in content. However, there were some differences. The results reported here are derived both from the interviews made and from examining the home pages.

Among all subjects 16 had included information about what they are working with, either what projects they are involved in or what areas they are working with, on their home page. The reasons reported by the subjects for including this kind of information were because they wanted other persons to know what they do, and to distribute information about projects. Those that did not include any such information said that they wanted to include it.

Information about how to be contacted, e.g., email, telephone number, postal address etc, was included by 21 of the subjects. In the interviews it emerged that this was included to give other persons the possibility to make contact, or because it was included on the page they used as a template. The one subject that had not included any contact information said that it was a mistake, and that it should have been included.

A picture of the author was included by eleven of the 21 subjects. According to the interviews they had a picture because they wanted the page to look personal, and because they found it useful to see a picture of a person they were going to meet but that they had not met before, and they therefore thought that others had similar preferences. The reason given for not including a picture was that those subjects did not have any picture that they considered good enough, or that those subjects only had one personal home page and if they should have a picture they wanted to have it on a page linked to from the top page. They wanted to avoid pictures on the first page because pictures take time to load.

Information about hobbies was only included by seven of the persons interviewed. This was said to be included because it was considered to make the page more personal. However, some subjects expressed that they did not want any personal information on the page either because they thought it did not belong there, as was also found by Bly et al. (1998), or because they did not want other people to know too much about their personal life.

Table 5.1, Table 5.2 and Figure 5.1 summarise what kind of information people included on their home page. All these categories of information can contribute to the four principles characterising a knowledge net (see Chapter 4).

	Contact info.	Proj.	Prev. proj.	Publ./ doc.	Edu-cation	Interest areas	Personal interests	Picture
dev (8)	8	7	1	2	1	4	3	4
inst (6)	6	5	2	4	2	3	2	4
uni (8)	7	6	0	3	3	8	2	5
total	21	18	3	9	6	15	7	13

Table 5.1: *The number of persons including each type of information on their home page.*

Two of the subjects from the dev group used their personal home page as a project

	Contact info.	Proj.	Prev. proj.	Publ./ doc.	Edu- cation	Interest areas	Personal interests	Picture
dev	100%	87%	12%	25%	12%	50%	37%	50%
inst	100%	83%	33%	67%	33%	50%	33%	67%
uni	87%	75%	0%	37%	37%	100%	25%	62%
total	95%	82%	14%	41%	27%	68%	32%	59%

Table 5.2: *The number of persons in percent including each type of information on their home page.*

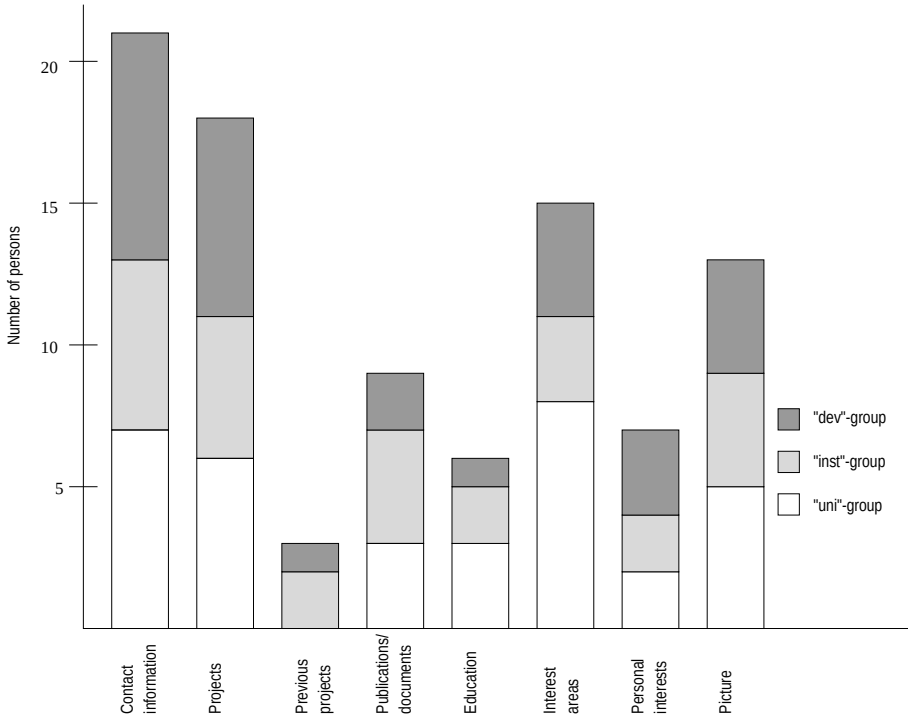

Figure 5.1: *Contact information was the most common type of information to include on the home page, but also information about projects was quite common.*

page, i.e., to distribute information about the project they were responsible for. This way of using personal home pages was not seen among the researchers.

Comparing these figures with the ones presented by Bly et al. shows a higher number of pages with a picture (39% in the study by Bly et al.compared to 59% in this study) and project information (75% in the study by Bly et al. compared to 82% in this study) included, but a lower number of pages with hobby information (or personal interests) (50% in the study by Bly et al. compared to 32% in this study) included. However, weak spots in these differences are the low number of subjects and the definition of the

categories used, i.e., the definition of personal interests used in this study may differ from the definition of hobby information used by Bly et al.

The use of personal home pages as a simple knowledge net

Most researchers (13) said in the interviews that other colleagues, known or unknown, could be interested in reading their home page. They thought that information about research projects, publications and contact information would interest such persons most. Also students were mentioned, by six of the researchers, as a possible group that could be interested in their home page. Students would, e.g., be interested in information about courses taught by the author of the page.

Six of the software developers, on the other hand, thought that colleagues within the group or the project could be interested in reading their home page. The information these persons would be interested in was mainly project information.

Half of the persons interviewed (11) had at least once been contacted by an unknown person who had read something on their home page. In one case it was a student asking about a thesis project, and in another it was a person who needed participants with a certain competence in a project. Also, some complained about what was on the home page. In one case the person had a picture that failed to load, and in another case it was something the person had written in a report, that was linked to from the home page, that the contacting person found to be wrong. In another case the contacting person was looking for some information on the subjects' home pages but did not find it and, therefore, contacted the owner of the page. It appears that having a personal home page on the Internet does not imply that you get a large amount of email from persons you do not know. On the other hand, the results seem to show that the home pages actually reach their target audience. However, people may not always be aware of where the sender of an email has found their address.

All persons interviewed said that they found the information on other persons' home pages useful. It was mainly contact information, publications, and project or research information that was found interesting. But some persons also said that it was one way to find out more about an unknown person that is to be contacted—then the person is not really unknown anymore. It appears that personal home pages facilitate contact between the author and other persons. This is also one of the main objectives of a knowledge net.

Do people keep their home page up-to-date?

When asking, in 1996, if the persons interviewed had made any major changes to their home page since it was created, it was found that only six persons had not made any changes. Ten persons said they had made major changes more than once. The changes made were adding text, adding links, changing the layout, adding a picture etc. The reasons reported for making these changes were that the page needed to be updated, new interesting links had been found, new projects had been started, more information was added which made a new layout necessary, etc.

When comparing each persons' home page from 1996, 1997 and 1998 it was found that most people had made some changes to their page during this time. However, only one person from the dev group had made changes to the home page, and they were minor changes. The changes mostly concerned the text written on the page but there had also been major changes in the layout.

Since only a few pages from the development organisation were investigated a second and third time it is difficult to draw any conclusions whether there are more or less changes in such an organisation than in the research organisations. However, when interviewing three persons in the development organisation in 1997 it was found that they were less motivated in keeping the page up-to-date because they were not certain about whether their home pages could be reached by others in the world wide organisation. In the follow-up study in 1998 it was found that the department to which the persons studied belong did not "support" personal home pages. The pages were not linked to from any other page in the organisation, but they could be reached from the world wide organisation. However, other departments within the organisation did support personal home pages by having links from organisational pages to home pages of people working in that department. The fact that the personal home page was not supported by the department can be an explanation to why the subjects from the dev group were less motivated in making changes to their home pages. The possibility for other persons to easily find the home page is obviously important.

Conclusions

The study has given strong indications to what should be included in a knowledge net. Contact information and projects or work that the author is involved in were information items that the subjects both found interesting on other persons' home pages and included on their own home page. The level of detail of the information is, of course, relevant. The more detailed the information is the better the references given will be. Therefore, a possibility to detect what function the person has in a certain project, and what parts of the project the person is working with would be an advantage.

Contact information is important in a knowledge net in order to know how to contact the person being referenced. This kind of information was included by almost all subjects.

One interesting result was that some of the subjects said that looking at an unknown persons' home page made them feel more familiar with that person. This is interesting because in a knowledge net people may be suggested persons that they do not know at all, and they may feel awkward to contact this person. However, if the persons referenced by the knowledge net are also presented in a way similar to how people present themselves on a home page, then this awkwardness may decrease.

The study shows that the subjects tend to update the information on their personal home page. This is an interesting finding because the willingness to keep personal information updated in a knowledge net is important.

It seems that one main group of intended readers of people's home pages are colleagues. The information on the home pages describes, in most cases, what "project"

the owner of the page is involved in and what areas s/he is interested in. The subjects presented what they had knowledge about rather than their opinions (although this may also be found on some pages). They did not seem to be afraid of presenting what they knew, wanted to do, and had been doing.

Also, that subjects found the information included on other persons' home pages of interest suggests that what is presented on personal home pages is of relevance for a knowledge net. This, together with the findings that some of the subjects had been contacted about their information on their home page, shows the interest for personal information. In this regard, many people may have, consciously or unconsciously, used personal home pages as a knowledge net.

The study shows that personal home pages can 1) provide knowledge about what information to include in a knowledge net, 2) serve as a source for this information, and 3) even be used as a simple version of a knowledge net.

The simple way of distributing and reaching information that the Web supplies is a valuable asset in an application such as a knowledge net. However, a knowledge net should be more sophisticated than a set of personal home pages in an organisation. One first improvement is to introduce a rating mechanism showing how much a person knows about a topic. This is one of the interesting questions in the next study.

Chapter 6

Study II: Letting individuals within an organisation describe their knowledge

Introduction

What kind of information do people include in a knowledge net? How can the information be entered? To avoid unreasonable workload on expert persons, how can a reference be described in order to find the "right" person and not the one that knows most? If the person that knows most is the one that is referred to then this person would receive a huge amount of questions. To get answers to these and other questions a prototype of a knowledge net was developed and tested.

One important design aspect of the prototype was to allow the users a large amount of freedom when choosing what topics to enter and how to describe their knowledge about the topics. The reason for this was to avoid restrictions on what information could be entered into the knowledge net. Any restrictions would reflect the opinions of the person conducting the study. For example, who is to decide what topics are suitable and useful to find in a knowledge net but potential users? Now the limitations were, to some extent, set by the users themselves.

The prototype was developed using the Web as a platform, and implemented using Java code together with Java Markup Language (JML) (Bogdan & Şandor 1997) and HTML. The Web was chosen as a platform because it is convenient to use for prototype development. It is easy to make fairly good user interfaces relatively fast. Also, to use forms seemed to be very useful when collecting information into a knowledge net. JML was chosen because it gives the possibility to separate the underlying structure from the user interface. JML also makes it possible to use Java objects and it can be used together with HTML code.

All data entered were stored under the UNIX file system using Java streams. To

use a regular database was unnecessary with the limited amount of data collected in the study.

Note that this study was only concerned with entering information in, or setting up, a knowledge net. There were no advanced search mechanisms or any advanced functions for presenting the information in the knowledge net prototype. Topics could be selected from a list of topics and the information about those topics could only be viewed in an alphabetically ordered list. Thus, it was not a prototype of a complete knowledge net.

Method

The study took place in two different organisations, one consisted of researchers and one consisted of software developers. Each part was divided into two sessions. During the first session the knowledge net was empty. Everyone was to enter new data into the knowledge net. During the second session the knowledge net consisted of all data entered during the first session. Now, everyone had the possibility to look at other persons' data, enter new data, e.g., rate or make relations between topics that was entered in the first session, and change or remove data they had entered before.

One kind of data to be entered was the name of topics together with a definition of the topic. At what level of detail the definition was to be made was decided by each person. Although we wanted to minimise the limitations on what kind of topics to enter the subjects were instructed to enter topics that they knew well or wanted to know more about.

Another kind of data to be entered was a rating of the user's knowledge about topics. Two different methods were used. In the first method people rated their knowledge using a scale between 1 and 7, where 1 corresponded to the situation that no one within the organisation knew less about the topic, and 7 corresponded to the situation that no one within the organisation knew more about the topic. This method, here referred to as the scale method, was chosen because it is simple and easy to understand, and it is easy to compare different persons' ratings. It also gives an indication of how well people relate their knowledge to the knowledge of other persons in the organisation.

In the second method people rated their knowledge about a topic by relating the knowledge to what they could do with it. Each person described an activity they could perform using the knowledge, an audience involved in this activity, and a setting in which the activity can take place. The activity, audience and setting were to be included, but not necessarily all three, in the following template:

> I know enough about *a topic* to do *an activity* involving *an audience* in *a setting*.

One idea behind this template (the idea was given by Professor John Bowers) is to describe the knowledge through an activity in which the knowledge is used. This method, here referred to as the activity method, was chosen because it is believed that it gives some context to the knowledge the person has. This could simplify the interpretation of a rating.

Both methods were to be used by the subjects when rating their knowledge about a topic. The knowledge the person had about a topic that was entered did not necessarily have to be rated.

A third kind of data to be entered was relations between different topics. To minimise the complexity of the prototype only one kind of relation was used; is_related_to. To enter relations makes it possible to search for related knowledge about a topic by following relations. Instead of having computed links, which may not always be correct, these links were explicitly created by the users themselves.

The subjects were instructed before the session started to think aloud when they used the prototype. This method was chosen because it can help to detect parts that people find difficult to use by making explicit how they think when they enter and look at the information in the knowledge net. In order to catch everything that was said during the study all subjects were tape recorded. The tapes from the sessions with the researchers have been fully transcribed.

It was also found to be important to see how people used the prototype, e.g., in what order they performed different actions, if the interface appeared to be easy to understand etc. Therefore, the researchers were also video recorded during both sessions of the part. The subject group of software developers were not video recorded because when studying the researchers I noticed that managing a video camera made me less focused on what was happening during the study. It was also difficult to get a good quality of the video (where all actions of the subjects were clear). Instead, the software developers were observed and notes were taken during both sessions of the part.

The knowledge net prototype was modified radically between the two parts in the study, and it was also modified a little between the sessions in each part (see below for more detail). These modifications were made according to the comments given by the users of the prototype in the different sessions.

Part I: Researchers

In this part 16 researchers (only 14 in the second session because two persons could not take part again) from a university used the prototype. All researchers belonged to the same research laboratory, but some also worked part-time in another laboratory. Only a few of the researchers worked together in projects.

In the first session of this part the researchers could enter new topics and rate their knowledge about the topics. In the second session they could also relate topics to other topics. The first session lasted for a maximum of 30 minutes and the second for a maximum of 45 minutes.

This part was completed a few months before the second part started.

The prototype

The first version of the prototype consisted of several linked Web pages. Each page was divided into two parts: the functions that could be performed to the left and the presentation of the results of the function chosen to the right.

Figure 6.1 shows what the Web page looked like when the researchers started a session. The text presented on the right part of the page is an introductory text about the study that was given as a start page. The functions presented on the left part of the page were *enter*, *change*, *delete*, and *view*. The functions could be performed on *topics*, *ratings*, and *relations*. A help function showing some examples was also available. The possibility to relate topics was not added until the second session of this part.

Figure 6.1: *Prototype I (second session): The page used when starting a session.*

When entering topics a definition of the topic was also to be given, see Figure 6.2. The instruction given for this task was "What topic are you interested in finding out more about or informing others about?". This would inspire the subjects to think of topics that could be suitable for a knowledge net and, at the same time, give a large amount of freedom to choose topics to enter. The instruction text was presented together with an empty text field with the label "Name of the topic", and an empty text area with the label "Definition of the topic". The topic was saved by pressing the button "Save". In the second session the possibility to go directly to the "rating" page or to the "add topic" page after saving the topic was added. This feature was asked for in the first session of this part.

Topics that had been entered into the knowledge net could also be viewed, see Figure 6.3. When viewing topics the user could choose between viewing all topics, viewing topics defined by a specific person or viewing topics beginning with a certain letter. The name of the topic, the definition and the person who entered the definition were presented in an alphabetically ordered list, sorted on topics, as a result of the operation.

When rating a topic people were instructed to use both the scale method and the activity method, see Figure 6.4. The topic to be rated was chosen from a list with all

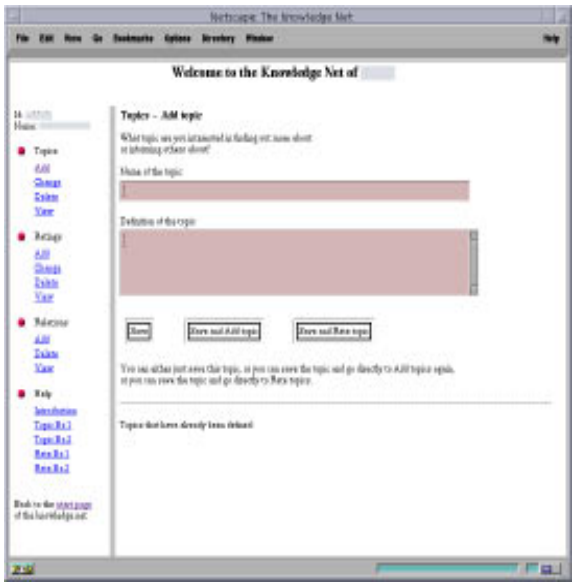

Figure 6.2: *Prototype I (second session): The page used when entering topics.*

Figure 6.3: *Prototype I (second session): The page used when viewing topics.*

topics that had not yet been rated by the user. The ratings were saved by pressing the "Save" button at the end of the page. In the second session the ratings other persons

had made of their knowledge of the topic were listed at the end of the page. It was also possible to choose between the activities, audiences and settings of the ratings of the specific topic that had been made by other persons. Also, the possibility to save the rating and go directly to the list with topics to choose from when rating topics was added. These features were asked for in the first session of this part.

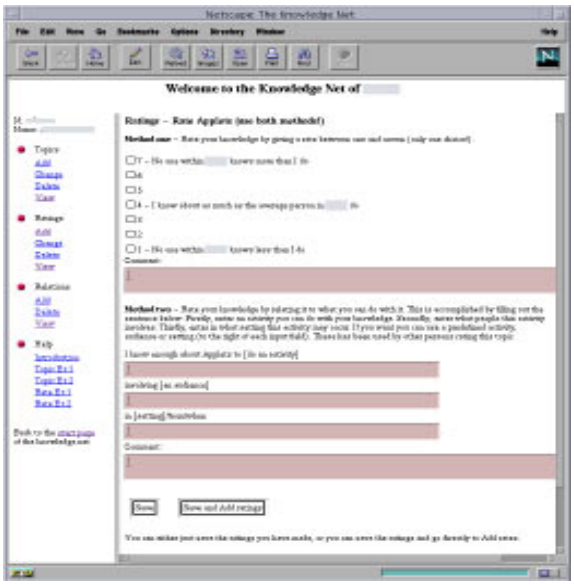

Figure 6.4: *Prototype I (second session): The page used when rating topics.*

The ratings could be viewed in a similar way as the topics, see Figure 6.5. The choice of viewing all ratings, ratings made by a specific person, or ratings of a specific topic could be made. The ratings were presented in a list, alphabetically ordered by the person who had made the rating, i.e., first all ratings made by "Don" were listed, then all ratings made by "Julia" were listed, and so on. The name of the topic and both ratings were presented in the list.

In the second session of this part the function to relate topics to other topics was added. This function was asked for in the first session. A relation was created by first selecting one topic from a list of all topics in the knowledge net, see Figure 6.6. In a second list of topics a selection of topics could be made of all topics to be related to the one that first was selected. The relations were saved by pressing an "Add" button. All relations that had been made earlier were listed at the end of the page.

As well as the topics and ratings, the relations could be viewed. People could choose between viewing all relations, or the relations of a specific topic, see Figure 6.7. All relations were presented in an alphabetically ordered list, sorted by topics. Since the function to relate topics was included in the second session of this part people could only view relations made by themselves.

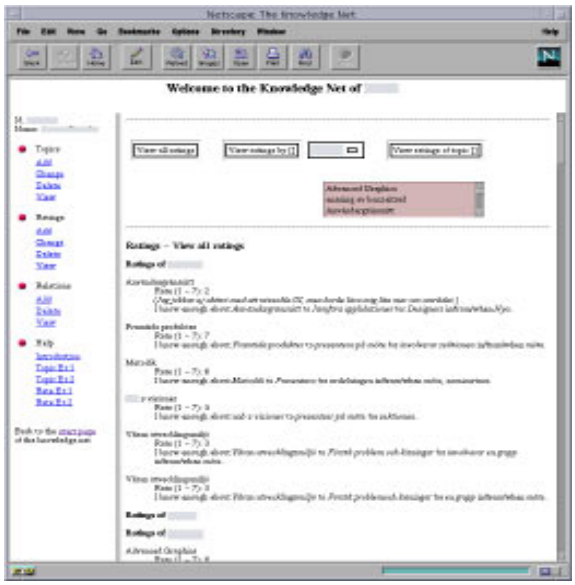

Figure 6.5: *Prototype I (second session): The page used when viewing ratings.*

Figure 6.6: *Prototype I (second session): The page used when entering relations.*

Part II: Software developers

In this second part eleven software developers (only ten in the second session because one person could not take part again) from a large development company used the proto-

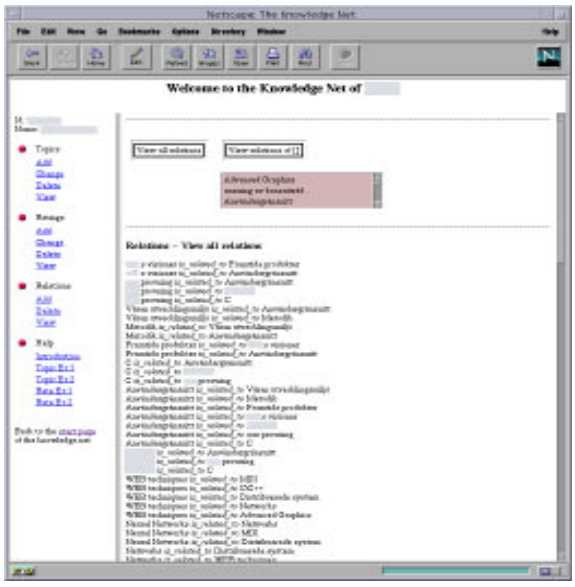

Figure 6.7: *Prototype I (second session): The page used when viewing relations.*

type. All eleven belonged to the same group, and nearly all worked in the same project, but on different parts. Thus, this group of people was more homogeneous than the group of researchers.

The software developers could enter new topics, rate their knowledge about the topics, and relate topics to other topics in both sessions. Both sessions lasted for a maximum of 30 minutes. The second session was shortened because it was apparent, from the first part with the researchers, that 45 minutes was too long.

The prototype

After the first part it was found that the interface of the prototype ought to be changed. The idea of grouping the functions around topics, ratings and relations did not work very well. People were confused about the differences of, e.g., viewing topics and viewing ratings. They wanted to see both definitions and ratings of a topic at the same time. People also expressed a wish to be able to operate directly on the list of topics presented when viewing topics and ratings. Thus, more focus was needed on the topics rather than on the functions.

Therefore, the interface of the prototype changed radically in this part, see Figure 6.8. This new version of the knowledge net prototype also consisted of linked Web pages divided into two parts. The left part consisted of a list of all topics that had been entered. A selection of what topics to show in the list was also possible to obtain using the buttons above the list. It was possible to list all topics, the topics defined by the

present user, the topics that had been rated by the present user, and the topics that had not yet been rated by the present user.

During this second part it was suggested that the possibility to create a “personal” sublist with topics of interest would be convenient. This would limit the amount of topics in the list to the ones that are of interest for the present user.

Below the list of topics a number of buttons representing different functions were located. In the first row there was one button for defining a topic, one button for rating a topic, and one button for relating topics.

The define button could be used either with no selection of topics or with a single selection, otherwise an error message was given in the right part of the page. If no selection was made then an enter topic part was presented in the right part of the page, and a new topic could be entered. It was the same enter topic part as was used in the first part, i.e., the same instructions, text field, and text area. If a single selection of a topic was made, and if that topic had already been defined by the user, then the same enter topic part was presented in the right part of the window but with the name and definition of the selected topic included in the text field and text area. This name and definition could then be changed.

The rating button could be used only with a single selection of a topic, otherwise an error message was given in the right part of the page. When rating a topic the same rating part as the one used in the first session of the first part (with the researchers), i.e., the same text fields and text areas but with slightly changed instructions, was presented in the right part of the page, see Figure 6.8. If the topic had already been rated by the user then these data were included in the text fields and text areas and could be changed.

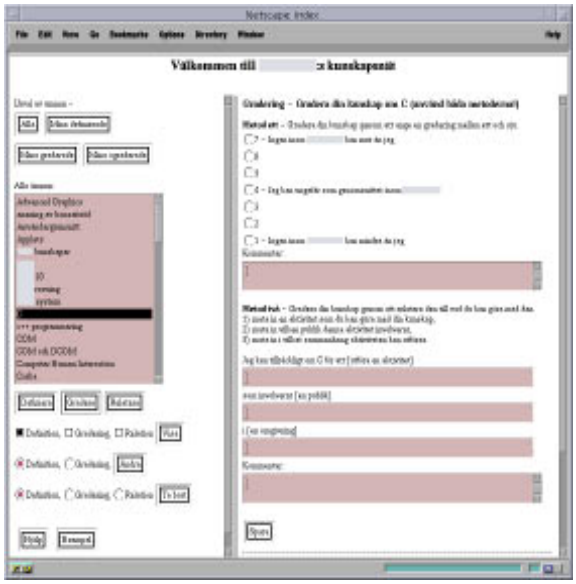

Figure 6.8: *Prototype II (second session): The page used when rating topics.*

The relation button could be used either with a single selection of a topic or with a multiple selection, otherwise an error message was given in the right part of the page. If a single selection was made and the relation button was pressed then a relation part was presented in the right part of the page, see Figure 6.9. A selection of one or more topics could be made from the list of topics in the right part of the page, and these topics were related to the one that was first selected. If a multiple selection from the topic list in the left part of the page had been made and the relation button was pressed then all these topics were related to each other. The changed procedure of relating topics were natural with the new version of the knowledge net prototype. Also, this "new" way of relating topics was asked for during the sessions with the researchers.

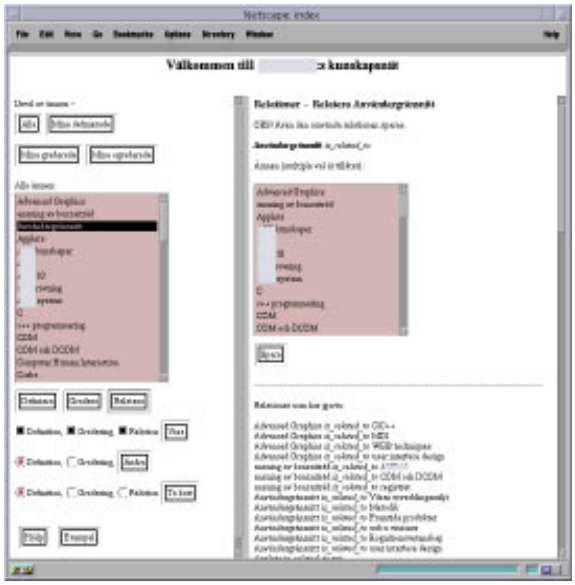

Figure 6.9: *Prototype II (second session): The page used when relating topics.*

Below the define, rate, and relate buttons the buttons for viewing the definitions, ratings and relations were found. A selection of viewing definition, rating, and relation of topics could be made. The data was presented in an alphabetically ordered list, sorted by topics, in the right part of the page, see Figure 6.10.

In the second session of this part it was also possible to operate directly on the topics listed to the right of the page when using the view function. By clicking on a "rate" link directly to the left of a topic name the user could go directly to the rating part for that certain topic. It was also possible to relate a topic in this way. If the person was the one who had defined a topic or entered a rating presented in the list then the choice to change these was also possible by clicking on links in the list.

Below the view buttons were the change buttons that could be performed on definitions and ratings. This worked in the same way as when selecting topics in the list

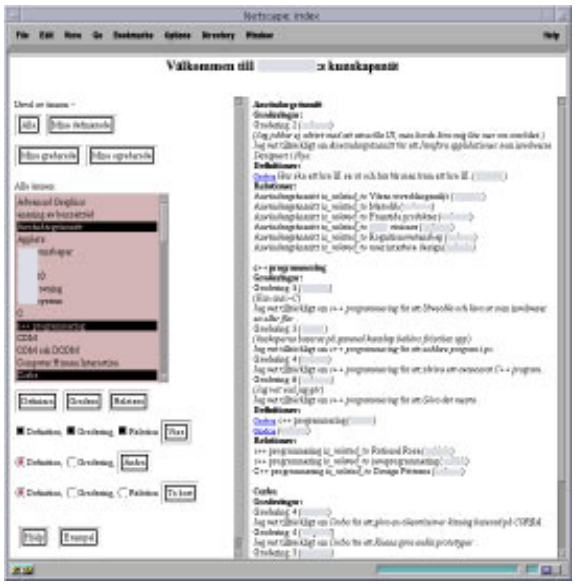

Figure 6.10: *Prototype II (second session): The page used when viewing definitions, ratings and relations.*

of topics and pressing the definition or rating button. Relations could not be changed because it is easier to first remove a relation and then enter the new one.

Below the change buttons were the remove buttons that could be performed on definitions, ratings and relations made by the user her- or himself. This function was used together with a selection of topics from the list of topics.

Below the remove buttons, at the end of the left part of the Web page, was a help function and a set of examples.

Results

Both parts of this study were completed according to the plans. Most subjects found the task interesting and they entered plenty of information. Some subjects were finished before the time was up, but still had entered plenty of information. Other subjects took longer time, and were interrupted when the time was up.

The number of topics entered, and ratings and relations made, varied among the subjects. For example, one subject only entered one topic and made one rating during the first session, while another subject entered eleven topics and made eleven ratings. Another example is the amount of relations made. One subject made 136 relations (including the converse relations that was created automatically), while another subject did none.

Topics

The results showed that the subjects tended to enter topics that were on a quite high level of abstraction. They rather entered topics like "virtual reality" than specific areas within virtual reality. It appears that these topics were natural to start with.

The topic definitions varied a great deal. If the name of the topic was written as an acronym then the acronym was simply written out:

> *HCI* – "Human Computer Interaction"

There were also longer definitions, especially among the researchers:

> *CORBA* – "Common Object Request Broker (CORBA) is a distribution effort made by OMG, an organisation of more than 600 companies and universities. The objective is to define a complete distribution package. CORBA is platform and programming language neutral, i.e. it enables you to easily communicate between for instance C++, Smalltalk, Java, Ada, etc."

Some topics were defined by describing where more information about the topic could be found:

> *Java* – "www.javasoft.com"

Other topics were defined using associations or synonyms:

> *Sociology* – "human, society, politics etc . . . "

Most topics entered were related to work, but there were some related to leisure time activities as well. This occurred more often among the researchers than among the software developers where only one hobby topic was entered (however, the subjects enjoyed reading about this topic and one other subject rated her/his knowledge about the topic in the second session). An example of a hobby topic is

> *Sarek* – "Sarek is a national park in Lapland, Sweden. Bad weather conditions, lots of rivers, mosquitos, great glaciers."

The number of topics added and defined per subject were around six among both the researchers and the software developers (the software developers used the prototype for a shorter time in the second session than the researchers), see Table 6.1. Not many topics were changed in the second session and only one was removed (this one has not been removed in the table because other subjects rated their knowledge about the topic in the second session of the study).

The subjects entered more topics during the first session, which is not surprising since there were no topics present at all in the first session of the study. Hence, the focus of the task shifted from entering topics and rating one's knowledge about one's own topics to rating one's knowledge about topics that other subjects had entered.

	total	average	st. dev.
Researchers			
session one (16)	86	5.4	2.3
session two (14)	16	1.1	1.1
in total	102	6.4	2.1
Software developers			
session one (11)	67	6.1	2.4
session two (10)	3	0.2	0.8
in total	70	6.3	2.7

Table 6.1: *The number of topics added in the knowledge net after each session.*

Some of the subjects said, when asked why they did not enter so many topics after the second session, that they did not have any more topics to add. This could be due to the "experimental" situation. The number of topics added by the software developers in the second session was lower than the amount added by the researchers. The researchers added on average 1.1 more topics per subject while the software developers added 0.2.

Performing *t*-tests (see Appendix B, Table B.1) on these results showed that the difference in the amount of entered topics in session one is not statistically significant at the level of 90%, while the difference in session two is at the level of 95%. An explanation for this difference can be the fact that the software developers used the prototype for a shorter time in the second session than the researchers did. However, there might also be a reason such that the software developers in their work are focused on fewer topics while the researchers are working with a larger number of topics

Ratings

The subjects found both rating methods somewhat difficult to use. In the scale method most of the ratings given were on the upper level of the scale, i.e., the subjects tended to use a sub-scale ranging from 4 to 7. This sub-scale does not have a middle alternative, which the subjects appeared to want since many found it difficult to choose between a 5 and a 6. Some of the researchers expressed that they wanted the possibility to give a "five and a half".

Some of the software developers, on the other hand, found it difficult to use the scale method because they felt they did not know the other persons in the group well enough. They had no idea of what some of the others in the group knew about specific topics. Some of the software developers also expressed that they were afraid of bragging when rating their knowledge using the scale method.

In the activity method the subjects found it especially difficult to enter a setting where the activity could take place. They also found it somewhat difficult to name an activity and an audience.

It also seems that the subjects found it difficult to interpret what knowledge people have about a topic using the scale method. This was, however, better achieved in the activity method where it was possible to get a picture of what it is the person knows about the topic.

When rating their knowledge about different topics the researchers wanted to be

able to specify what parts of the topic they had knowledge about. On the other hand, they did not choose to specify the topics on a lower level. Instead, they did as in the example below where the subject has made a comment excluding some aspects of the topic.

> *Scale method*: "5 (Excluding knowledge about technical aspects on e-mail)"
>
> *Activity method*: "I know enough about E-mail to give a lecture for undergraduate and graduate students."

To see how well the two rating methods correlate, 13 of the researchers were given a list of all topics, together with the ratings made with the activity method, that more than one person had rated their knowledge about. The list consisted of 29 topics and two to five activity ratings for each topic. For each topic each person ordered the ratings by "who-knows-most" about this topic. It turned out that nine of the orderings made by the subjects correlated exactly to the order of the ratings made with the scale method, and that eleven correlated closely. Accordingly, nine of the ordered topics did not correlate at all with the ratings made with the scale method.

The ratings made by person A and person B in the example below is an illustration of the correlation between the two rating methods. All researchers classified person A's activity as having more knowledge about the topic than person B's activity.

> Person A:
>
> *Scale method*: 5
>
> *Activity method*: "I know enough about HCI to give a small course involving students and software designers from anywhere."
>
> Person B:
>
> *Scale method*: 4
>
> *Activity method*: "I know enough about HCI to I have taken two graduate courses: Methodologies of HCI, Design for HMI."

It was also quite common that the subjects, especially the researchers, entered more than one activity, audience and setting:

> *Activity method*: "I know enough about object oriented programming to program, analyse, design, teach, write (about), find and describe a design pattern involving programmers, customers, students, readers in fabrication of software, classes, seminars."

Sometimes the subjects found it difficult to rate topics they did not know much about using the activity method. However, the example below shows that there were creative ways of doing this.

> *Scale method*: 1 "(Perhaps there are others who don't know any knitting)"
>
> *Activity method*: "I know enough about knitting to recognise a knitting person. (I never got around to learning knitting in school.)"

The comment fields were frequently used. In the example below the subject used it to point out that the 6 on the scale soon would be a 7.

> *Scale method*: "6 (In the future I want to rate the topic with 7, but at the moment I don't have that much experience. This easter I will go to Sarek for my first winter tour in that area.)"
>
> *Activity method*: "I know enough about Sarek to guide some people with a good life insurance, if they don't have that, then we go to Kebnekaise instead."

	total	average	st. dev.	method1	method2
Researchers (16)					
session one	81	5.1	2.3	81	78
session two	67	4.8	3.8	67	53
in total	146	9.1	5.6	146	131
Software developers (11)					
session one	59	5.4	2.9	59	55
session two	37	3.7	2.9	37	31
in total	94	8.6	4.6	94	84

Table 6.2: *The number of ratings in the knowledge net after each session. (Some ratings were changed in session 2, and that the total number may not be equal to the number of ratings entered in session one plus the ones entered in session two.)*

The number of ratings entered per person were around five among both the researchers and the software developers in session one and around four in session two, see Table 6.2. The difference between the two groups was not statistically significant at the 90% level (see Appendix B, Table B.2). Together with the low number of new topics added in the second session this large amount of ratings made in the second sessions indicates that the subjects found many of the topics added by other persons to be of interest. Not many ratings were changed in the second session, but some were. This is the reason why the total number of ratings in Table 6.2 are not always equal to the total amount of ratings entered in the two sessions.

The scale method was used all the time by all subjects, but the activity method was sometimes left out. Mostly, this was when the subjects rated their knowledge about the topic on the lower part of the scale (i.e., implying that they had limited knowledge about the topic).

The different ratings the subjects gave using the scale method were concentrated around 5 and 6 among the researchers and around 3, 4, 5 and 6 among the software developers, see Table 6.3 and Table 6.4. When calculating the average rating given by the researchers and the software developers in each session it showed that the researchers gave a higher average than the software developers in both sessions (see the last line in Table 6.3 and Table 6.4). Even though this difference is not statistically significant at the 90% level (see Appendix B, Table B.3) it is tempting to interprete it as an effect of that the researchers are more specialised within their area of work and, therefore, also have more knowledge about the topics they entered.

Rating	Session 1		Session 2		Total	
	Number	%	Number	%	Number	%
1	1	1%	2	3%	3	2%
2	3	4%	1	2%	4	3%
3	4	5%	2	3%	6	4%
4	13	16%	15	22%	27	16%
5	27	33%	27	40%	53	36%
6	29	36%	17	26%	46	32%
7	4	5%	3	5%	7	5%
average	5.12		4.92			

Table 6.3: *The number of topics given a certain rating by the researchers in the two sessions. 1 corresponds to "no one knows less than I do", and 7 corresponds to "no one knows more than I do". (Some ratings were changed in the session 2 and that the total number may not be equal to the ratings added in the two sessions.)*

Rating	Session 1		Session 2		Total	
	Number	%	Number	%	Number	%
1	0	0%	2	6%	2	2%
2	6	10%	2	6%	8	9%
3	8	14%	8	23%	16	17%
4	11	19%	6	17%	17	18%
5	11	19%	13	34%	23	24%
6	17	29%	5	11%	21	22%
7	6	10%	1	3%	7	7%
average	4.73		4.53			

Table 6.4: *The number of topics given a certain rating by the software developers in the two sessions. 1 corresponds to "no one knows less than I do", and 7 corresponds to "no one knows more than I do". (Some ratings were changed in the session 2 and that the total number may not be equal to the ratings added in the two sessions.)*

In the first session the subjects gave topics a higher rating than in the second session. There can be many reasons for this. One can be that in the first session people entered topics they know well, and in the second session they found additional topics, entered by someone else, that they did not know as well as the ones they entered themselves. Another reason can be that in the second session the subjects could compare their own knowledge to that of others.

Since the subjects related their knowledge about a topic to how much other persons within the organisation know about the topic this upper sub-scale effect shows that the topics entered are usually the ones where the subjects can answer questions. That they provide this kind of topics indicates that the subjects are willing to share their knowledge with others.

When examining each subject's home page (only seven of the eleven software developers did have a home page) it was found that about 30% of the topics rated could be found on the home page. However, for some subjects none, or very few, of the topics they had rated their knowledge about were found on their home page, while for others

many of the topics were found on their home page. Of course, the kind of information represented on each person's home page varied. On some home pages there were only published reports, the name of the projects the person were working with and contact information. On others there was also a longer description of what the person was interested in and working with. Also, one person decided not to enter topics that the other persons clearly would be aware of that s/he would know well. The person expressed, when asked after the study, that s/he found it more interesting for others in the group to know what s/he knew about other topics.

Relations

About 75% of all topics were related to another topic. Some topics were only related to one or two other topics, while others were related to several topics (e.g., the topic Java was related to 12 other topics among the group of researchers).

Some relations were created because one topic could be used within the area of another topic:

Interviews *is_related_to* HCI

Other relations given were created between topics of the same category, e.g., geographical places or leisure activities:

Sarek *is_related_to* Pakistan

Sports *is_related_to* Cooking

Such relations may not be useful in a knowledge net application. However, people were instructed to make any relations they wanted and, apparently, there was a need to group topics according to such criteria.

Also, topics were related if they were meant to be the same but the knowledge net prototype distinguished between them because they were spelt differently:

Object Oriented Programming *is_related_to* Object-oriented programming

This type of relation can be avoided with better mechanisms for distinguishing between different spellings of the same topic. At this time the knowledge net prototype only treated capitals and small letters as equal.

The number of relations made (each relation made resulted in two relations, the one made and its converse) per subject was around 35 (in the second session) among the researchers and around 11 among the software developers, see Table 6.5. The total amount of relations added was higher among the researchers than among the software developers. Also, the total amount of topics that were related were larger among the researchers than among the software developers. However, it is not possible to draw any conclusions about these differences because the relation procedure was used differently by the two groups. The researchers used it only in session two, while the software developers used in both sessions. Also, none of these differences turned out to be statistically significant at the 90% level (see Appendix B, Table B.4 and Table B.5).

Some persons in both groups related topics and some did not. However, some persons among the researchers entered many relations (the one that added most relations among the researchers entered 136, while the one that added most relations among the software developers entered 30). This variability is interesting. Either the task was not clear enough to all persons, or people have different criteria for what counts as a relation. Some subjects may have been very liberal when creating relations, while others may have been more conservative.

	number of relations	average	st. dev.	number of topics related	average	st. dev.
Researchers						
session one	-	-		-	-	
session two	486	34.7	(37.7)	75	17.6	(13.7)
Software Developers						
session one	70	6.4	(7.2)	35	3.4	(2.4)
session two	52	5.2	(7.6)	25	2.5	(3.9)
in total	122	11.1	(9.3)	47	5.6	(4.1)

Table 6.5: *The number of relations added by the subjects in the two parts in the first and second session of the study.*

Conclusions

A large number of topics were entered into the knowledge net together with a large amount of ratings of the knowledge people had about the topics. None of the subjects expressed that they had knowledge that they did not want to share with others. Also, there were no indications that people felt uncomfortable in making their knowledge public within the organisation. This suggests that people are willing to share their knowledge with others in the organisation.

The topics entered and rated were not only topics of interest for the current "projects" the subjects were working with. Also topics that the subjects found interesting, or wanted to learn more about, were entered and the subjects rated their knowledge about these topics as well. This suggests that people are also willing to share knowledge they have but that they are not "expected" to share. This also suggests that support for creating informal interests group about different topics can be an interesting feature in a knowledge net. If such interest groups are organisation wide then the knowledge about such a topic will be used to a greater extent in the organisation.

Two simple methods were used to give a rating of peoples' knowledge about different topics. The scale method appears not to work very well because it is relative to other persons' knowledge and it does not say anything about which person is most suitable to ask. On the other hand, since it tells how much a person knows, it is measurable and it can help make a first filtering of the ratings. Another disadvantage with the scale method is that it does not have a reference point, which explains the large variability in the ratings made by the researchers and the software developers. However, supplying a reference point for a topic, e.g., showing the average knowledge about the topic in the

group, is not possible without having knowledge about all persons' knowledge about the topic. Not even the subjects themselves knew how much the other persons in the group knew about the topics. One other finding that also makes it difficult to use the scale method was that a few persons felt that they were bragging when comparing their knowledge with others in the group. No one expressed this when they were using the activity method.

Using the activity method appears to be useful. However, how the activity, audience and setting are to be entered is critical for the understanding of the method. Many of the subjects, in both parts, found it especially difficult to define a setting for the activity. Also to define an audience for the activity was found to be somewhat difficult. Reasons for this could be that the instructions were insufficient or that the subjects, in some cases, entered an audience or setting together with the activity. On the other hand, the method was very useful when interpreting the ratings. It indicated what it is the person knows, not only that the person knows more or less. This is most likely important when trying to find the person most suitable to give as a reference. Still, this method has to be complemented with a measurable method.

The number of relations made by the subjects indicates a reasonable amount of related topics in a "real" knowledge net. For the relations to be useful in a knowledge net a more conservative use is to be preferred because if there are too many relations then the relations are of little help.

One aspect that needs to be considered is how the results from this study have been affected by the experimental situation. Even though the subjects were aware of the fact that everybody else in the group would see the information that was entered in the first session of the study they may still have been more "relaxed" in their use of the prototype than they would have been if it would have been a "real situation". They never had to answer to the knowledge they claimed that they had.

The positive reaction to the knowledge net prototype suggests further work on developing additional functions, such as a search mechanism and a graphical presentation of people with knowledge of a certain topic. The last version of the prototype, the one used by the software developers in the second session, can still be improved, but the overall design is probably suitable for the tasks tested in the study. To let potential users, as subjects in a study, use and influence the design of the knowledge net prototype has shown to be useful and will also be used in future work.

Chapter 7

Study III: Studying knowledge sharing between colleagues in an organisation

Introduction

A study using ethnographic methods has been performed in order to find out how people in an organisation find the right person to ask or discuss problems with. Also, how aware people are of each others' activities and who they choose to ask about different matters were questions to be answered by the study.

A group of eleven software developers (the same as one of the groups used in Study II, see Chapter 6), among which one is manager of the group and another is an official expert resource for the whole department to which the group belongs, has been observed and interviewed in their workplace. The group is part of a larger development company that, in turn, is part of a world wide organisation.

The world wide organisation has an intranet structure that can be reached from all companies that are part of the organisation. The company where this study took place, and some of the other companies part of the world wide organisation, also have a local intranet structure. The local intranets appeared to be reachable from all companies part of the world wide organisation.

The local company is divided into different parts, which in turn consists of several departments. Each department consists of several groups. Each group consists of 10 to 20 persons.

Related work

A very interesting and similar but larger study has been made by McDonald & Ackerman (1998). In their study of a software firm two mechanisms for finding persons to ask

were identified. The first was expertise identification that was performed to know what kind of skills and knowledge other individuals had. The second was expertise selection that was performed to select the individual with the required expertise.

McDonald & Ackerman found that for the expertise identification procedure, people used historical or archival data that the organisation maintained, and that they turned to people that were considered to have the role of an "expertise concierge". They also found that people felt it difficult to articulate how they know who had a certain kind of knowledge. In the expertise selection procedure the authors found that people used different organisational criteria. One criterion was that people preferred to keep the question local. If this did not work then they crossed departmental boundaries. People also turned to the "expertise concierge" of the organisation. Further, they found that people selected other persons on the basis of workload, both regular workload and workload over time. Also how people share their expertise, how they perform in this matter, influenced if they were selected or not.

In a number of field studies Salvador & Bly (1997) found that people rely on interactions with other people and information sources to perform their work. The authors call patterns of such interactions "constellations as sets of people and information sources connected to a specific individual, which cross organisational and corporate boundaries and which exist in a particular work context". A constellation differs from a personal network in that "the people and information sources provide value in an immediate and specific context". A constellation also differs from a "team" in that "constellation members may not know each other and only contingently share common goals".

Salvador & Bly also describe one case of non-reciprocity in their interviews. Person A holds some persons in his constellation. However, these persons do not consider person A to be part of their constellation. Therefore, "the value relationship is one way". The authors also describe one case of reciprocity. From the beginning it was C who needed to discuss certain matters with B. However, over time B and C developed a relationship around the particular context. From the beginning B was part of C's constellation, but B's constellation evolved because s/he also gained from the contact with C.

The reciprocity of sharing knowledge is also something that is relevant for a knowledge net and that is touched upon in this study.

Method

In order to study people in their real work situation ethnographic methods were used. Ethnography can be considered to be a combination of various research approaches and analytic frameworks (cf. Hughes, Rodden & Rouncefield 1994, and Bowers 1995). Hughes, Rodden & Rouncefield also consider the main characteristic of an ethnographic approach to be that the researcher "unobtrusively" observes a work place by being present among the workers, and that it results in a large amount of descriptions of circumstances, practices, conversations and activities that reflect everyday work settings.

Ethnographic methods were chosen because the procedure performed to find other persons to ask about different matters was believed to be difficult to capture using in-

terviews only. To observe people in action would give additional information on how people act when they look for other persons.

All persons in the group studied, and one other key person, were initially interviewed. Questions about their background, their routines at work, their tasks at work, who they collaborate with and why, and how they search for information when they do not know where to find it were asked. The group, and also to some extent persons from the other groups of the department, was then observed during a period of seven weeks, and many activities and "conversations" were documented. The interviews were tape recorded and notes were taken during the observations. A number of different meetings were also attended during this period, and notes were taken during these. All email sent to the group as a whole during this time was collected, as were a number of Web pages.

Results

The department consists of four groups, which will be called group A, B, C and D. Group A was the group studied. People in group B collaborated very closely with people from group A, and they were located in the same corridor. People in group C collaborated in some activities with people from group A, and they were located in a corridor in parallel with group A's. Group D, on the other hand, had almost no collaboration activities with people from group A. They were also located in a corridor in parallel with group A's. How the groups were located, where the entrance was located, where the canteen was located, and where the stairs were located, is shown in Figure 7.1.

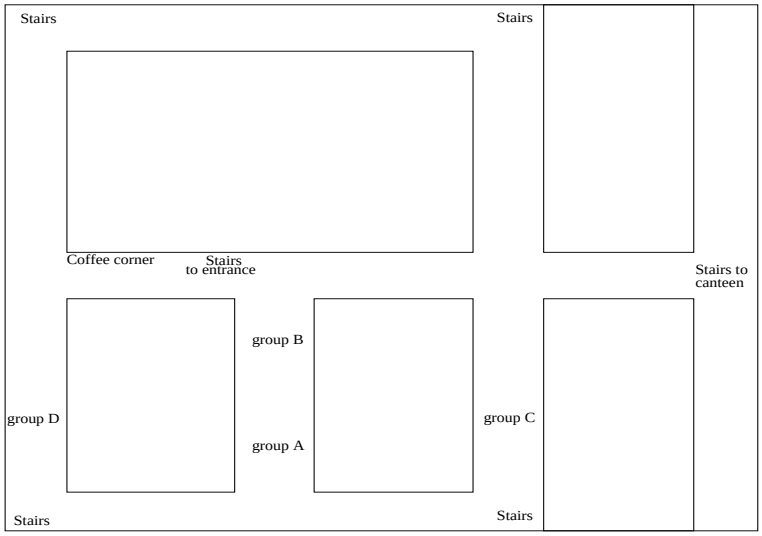

Figure 7.1: *A map of the building where group A was located.*

Walking around

The corridor in which group A and group B were located was full of activity during the day. People were walking in the corridor when they were coming or leaving for the day, going to meetings, going for lunch or a coffee break, or when they were looking for other persons.

Much of the walking around was initiated by people seeking another, specific person. In most cases it was persons working with the same applications but in other cases it was persons working with the same kind of software platform or programming language. There were a number of situations when this was observed, but some examples are:

> *Allen walks to Ambrosius' room and asks who in the project is travelling to the city of Btown tomorrow.*

> *Annie and Ambrosius walk together down the corridor of group A and group B. They are talking about setting a time for a meeting. They both walk into Ambrosius' room and Ambrosius checks his calendar. They talk some more and Ambrosius says "isn't Erland in his room or . . . ". Annie walks down the corridor to Erland's room and knocks on the door. After a while Ambrosius follows.*

> *I am walking from my room to the printer room to get a print out. Arthur is there making some copies of a paper. Arthur asks me some things that he has been thinking about that concern my study.*

What he wanted to tell me was not important enough to look me up and tell me, but when he saw me he took the opportunity to do so.

> *Ambrosius is arriving. In the corridor he meets another person, X, and Ambrosius asks "is the meeting now in the morning or . . . ".*

Ambrosius happened to see X when he came walking in the corridor and took the opportunity to ask. He did not explicitly walk around looking for an answer.

At other moments people were walking up or down the corridor to find somebody, not any specific person, to ask. This also occurred several times and some situations reported in interviews are:

> *Andy wanted someone to discuss some ideas with and therefore went out in the corridor to ask somebody. Anna happened to be in her room, and Andy discussed his ideas with her.*

> *Andy wanted to ask someone about how some specific software tools worked. Andy started walking around in the corridor and asked people who were sitting in their rooms. However, no one could help him, except for Aston who had some ideas that were of help.*

This is a clear illustration of the importance of locality as a selection criteria when looking for an expert.

> *Anna said that if she needs somebody to ask then she usually walks up the corridor and asks whoever she meets or whoever is in their room.*

Wandering around in the corridor could also lead to meeting someone that had been absent for some time. An observation was made of the following situation:

> *David from group D walks down the corridor of group A and group B. David stops outside Buzz's room and starts talking to Buzz. Buzz has been working in group B for a year but he has been away for several weeks. Another person, X, walks down the corridor, says hello to Buzz, and continues walking down the corridor. David leaves after a while. X comes back and, this time, starts talking to Buzz.*

Neither of David and X were explicitly looking for Buzz, but they stopped for a chat when they saw that he was in his room.

Another situation that occurred quite often was when two people were talking to each other and a third part came in and started talking to one or both of the persons. The three observations below illustrates this situation:

> *Axel comes into my room and starts talking. When Axel is about ready to go back to his room Assar (who I share my room with) starts talking to him.*

In this situation Assar noticed when Axel and I were finished and took the opportunity to discuss some matters with Axel.

> *Aston and Ambrosius are in Ambrosius' room discussing an application. Doris and another person, X, stop in the corridor outside Ambrosius' room. They are talking to each other. After a moment Aston walks out of Ambrosius' room, stops and starts talking to Doris and X.*

> *Allen is in Aston's room asking something. Aston says something about printing out, and they both walk down to the printer room. Aston comes back and Allen follows behind. Bonnie stops in the corridor outside Aston's room and asks Allen a question. Allen answers and then walks into Aston's room.*

Building contacts

When a person is new in the group s/he is given a contact person. This person helps the new employee with all kinds of questions or matters. Assar was new to the group while I was doing the study, and Alvin was his contact person. On Assar's first day at work Alvin walked around with him in the corridor and presented him to other persons of group A and B. Alvin also supplied Assar with contacts to talk to about different matters.

Arvid was also new to the group while I was doing the study. The manager, among others, supplied Arvid with a list of names of persons to be contacted concerning the task Arvid was going to work with.

The managers in the department are working on making their contact net broader. The manager of one of the groups said that the manager of one of the other groups has many contacts outside the organisation, and that the manager of yet one of the other groups also has quite a large contact net. The manager also expressed that s/he had to work much more on broadening her/his contact net.

Reciprocity

People in the group seem to always be open for questions. No situation was observed to testify against this assumption. The following observations show that people were open for questions:

> *One person, X, walks up the corridor and looks into Aston's room. X notices that Aston's room is empty and walks down to Anna's room. There X stops and says "hello hacker, do you have time for some questions". X and Anna start talking to each other.*

> *Eve (all persons with names beginning with an E did not belong to any of the four groups, they were, e.g., administrative personnel) has some trouble using the computer and walks to Erling's room. Erling asks her some questions about the problem and Eve answers. Erling follows Eve to her room and helps her with the problem.*

> *Eric walks down the corridor and stops at Aston's room. Aston says hello and Eric asks if Aston can join him in a meeting. Aston says "that is ok, when do we go". Eric answers "in about quarter of an hour", and then adds that it would be convenient if Aston can show some overhead slides.*

Of course, it can be the case that Aston knew that this was a very important meeting and that he therefore prioritised to go, or that he did not have other things that he had to work with for the moment. Still, it shows that people working within the same project are prepared to help each other. Whether or not people expect to be helped every time they look for other persons to ask is difficult to say. Nothing in the observations or interviews indicated such expectations.

Another example of the openness for questions is Aston who was a contact person for Assar that was recently employed. Assar frequently asked Aston questions and no observations were made of Aston not being open to answering the questions.

The corridor and "concentrating" knowledge

Many people in the group used small white boards outside their office to write messages on, and some put a post-it note on the door. Examples of messages people wrote are:

> *"On leave, back on Tuesday"*

"At home today"

"Lunch"

"Meetings Tuesday, at ACompany on Wednesday and maybe Thursday, mobile number is 000"

"Sick"

"At a course 3rd–5th, back on 8th"

"At company X"

"Away Tuesday–Thursday"

"At home with sick children"

"In ATown for two days, on leave next week, Bonnie is stand-in"

The awareness of when people are arriving and leaving, when they are going for lunch or for meetings etc is also a kind of knowledge that resides in the corridor. For a person located high up in the corridor (where people walked by arriving or leaving for the day, going for lunch etc) this awareness was stronger than for a person sitting far down. On the other hand, a person sitting far down in the corridor passes the other persons rooms and then becomes aware, by seeing an empty room, a closed door or a message on the white board, of if people are present or not.

One day the following was observed:

Ambrosius comes out in the corridor and asks if anybody is going for lunch. A group of people leaves for lunch and the corridor is now nearly empty. After about five minutes Arvid (sitting in a room located at the far end of the corridor) also comes to the canteen. He says that he did not notice the other persons leaving, and nobody had been thinking of passing his room (which is not on the way to the canteen).

Usually one person asked out loud in the corridor if anybody wanted to go for lunch. It very seldom happened that people walked down the corridor to see if anybody was still there. Walking up the corridor to the canteen, persons still in their room were asked to join for lunch.

There seemed to be two lunch groups, one that went earlier and one that went later. The one that went later also seemed to pick up the ones that were still in the corridor and that they were passing. The groups consisted of about the same persons every day.

Two examples of how movements in the corridor support awareness of the whereabouts of people are:

Several people are walking down the corridor and turn left in the adjacent corridor. It is clear that they are going to a meeting. Abe is standing in the corridor looking at the crowd walking down the corridor. I ask Abe if he knows where they are going. He says no but that he is also wondering.

Several people are walking down the corridor. It is clear that there has been a meeting and that it is now over.

The rhythms of a working day are more visible to people located where people usually are passing when they arrive, leave, go for lunch, go to meetings etc. When people are located in a corridor rather than in an "open plan office" there will always be an unequal distribution of awareness among people.

On various occasions people were using their room as a meeting room. In some cases the meeting was informal and only concerned two or three persons, discussing and examining an application. In other cases it was more formal meetings of a project with several people involved. Then a larger room with a small table was needed, e.g., the room of a manager. An example observed was:

Ambrosius, Annie, Anna, Arthur, and two others are having a meeting in one of these persons' room (which is a larger size room including a small table).

However, sometimes, not very often though, people closed their door for privacy. In one case a person was talking on the phone. To be able to close the door for privacy or for a meeting is one advantage with the corridor office rather than the open plan offices.

Contacting former personnel and personnel at former workplaces

People do keep in contact with former personnel and personnel at former workplaces. They also use these contacts when needed. An informal contact can be used to facilitate a formal approach:

Barb says in the interview that she has, on several occasions, arranged visits to an environment where the applications developed are used. She has former colleagues working there and uses her contacts to arrange these visits.

People also contacted persons they knew had shared or experienced a similar problem. This person was not necessarily a former colleague, but was not necessarily part of the organisation either:

Annie says in the interview that she knew that a problem she had, also had been experienced by a person at a different company, not part of the world wide organisation. She contacted this person who gave her some advice on the topic.

Using former colleagues, or even any colleagues, is also an alternative to the use of formal approaches if they are faster and less demanding. One example is the local help desk for computers which was regarded as slow and sometimes not competent enough. Usually a question took two to three days to be answered, if answered at all. Instead, some people had a former colleague working at the help desk office, and they contacted that person directly. Other alternatives was to contact a person from a former workplace that they knew was familiar with the computer environment.

Individuals with broader networks

If, in a group of people, one person has a broader contact net this is usually known by the others in the group. Examples from the interviews are:

> *Ally says in the interview "Chris is a person with much knowledge about other projects and it is he who meets other people".*

> *Andy says in the interview that Conrad knows most about similar applications developed outside the organisation.*

Except for the manager there was also an expert in the group who had a broad network of contacts. This person's main task was to support people in, mainly, the department within her/his expert area. This person was part of a network in her/his field of expertise where people from all over the world wide organisation participated.

The managers provided contacts to people when needed. This could, for example, be the case when entering a new project. In these cases usually the managers from two or three of the groups, and sometimes also the manager of the department, were involved. One of their tasks was to provide the software designers with contacts of whom to talk to regarding related projects etc.

Contacts were also given when an issue arose:

> *Betty says in the interview that at one time during the project her manager gave her the name of a person that could be contacted regarding an issue that had occurred.*

Another example is, as mentioned earlier, that new employees were supplied with name of persons to contact regarding the task they were going to work with.

The order in which persons are contacted

People tend to first ask a person nearby, and if that person does not know then they ask another. If one shares a room with a person then it is easier to ask that person first:

> *At one moment Assar, who I share my room with, asks me if I know about any tools for developing prototypes. I cannot help him, but in turn, I call across the corridor to Ambrosius, who I know may have this knowledge, and ask him if he knows.*

The order in which persons are asked also seems to depend on the seniority and accessibility of the person to ask. One example given in one of the interviews is:

> *Anya first talks to Bonzo and Arthur, whom she is working with, if she has any problems with Java. However, if they can not solve the problem then they go to Betty. They also know that Betty, in turn can ask Jimbo. Jimbo is considered to be the one with most knowledge about the Java code used.*

Another example is Ambrosius who often discusses Java problems with other persons from group A and group B who works with Java. If they cannot solve the problem then he searches for a solution on the Internet, e.g., Web sites and News, or contacts another, "external" person who might know.

Barriers to sharing information

Although people were willing to share their knowledge with others there also exist some barriers. In most cases, time was such a barrier. People working in one department may not prioritise spending time discussing problems or ideas with people from another department if it is not directly relevant to their work. One example from one of the interviews is:

> *Alvin needed a person to discuss some specific parts of his work with. Another person gave him the name of a person working in another department. This person was not the right person to discuss the problem with but he, in turn, could direct Alvin to yet another person, Tim. Tim on the other hand was very busy and his managers did not want him to spend time discussing problems with Alvin. Some of the managers in Alvin's department had to negotiate with Tim's managers and were eventually able to arrange a meeting where Alvin and Tim could meet and discuss Alvin's problem.*

However, time barriers can also exist within the same department. At one time Jimbo had very much to do in order to make a delivery. The project manager then said to everyone in the department that Jimbo was not to be disturbed.

Another example is when Anna started to work with a software application that another person was working with. This other person was leaving the organisation, but s/he was too busy so Anna did not get much opportunity to ask about the parts of the application that the other person was working with.

Another barrier to spontaneous sharing of knowledge is money. How is the time spent on discussing a problem to be paid if the person sharing his or her knowledge is not from the same department, or company, but in the same world wide organisation?

As an example, Axel had been asked by a former colleague, now working in another company within the organisation, to give a seminar about what he knew about a specific topic. The time Axel spent on this had to be paid for by the other company.

However, even though people knew that some persons had a high work load they still valued their competence and contacted them for a discussion about certain issues. Alvin said that Jimbo usually has very much to do but he still intended to contact him and use him as a discussion partner regarding certain issues.

A third barrier that may be a problem is security. During a meeting people were discussing the holiday schedule. They were not allowed to have the holiday schedule on the intranet, instead they used a paper version on the wall in the corridor. The reason for this was, according to one person in the group, to avoid having un-welcome visitors at home during the holiday.

Conclusions

The persons in the group studied were all familiar with each other to some extent, but the ones working with the same application or in the same project knew each other better. It can be an advantage to make use of a knowledge net as informal as when asking a close colleague. To be aware of who the person referenced is, where he or she is located and what s/he is working with etc, can in this regard be important.

In many cases the subjects solved their problems within the group. The need to ask someone outside the group or project seemed to arise less often in the every day work. On the other hand, most people interviewed were not aware of whether similar activities were or had been performed, or if other people were working with similar software tools or programming languages, in other departments in the local company, or in other companies within the world wide organisation. The study showed that most people in the group had a rather limited personal contact net. They knew who to ask in the corridor or among people participating in the same project. However, only a few of the persons studied had contacts outside the department. This is where a knowledge net would be useful—to find other persons with experience within the same area of work.

Different barriers for sharing information across local companies that were part of the same world wide organisation were identified. It was sometimes a question of time and money when a person was doing something not part of the present project. These barriers need to be taken under consideration when designing a knowledge net. Another barrier that occurred was security.

There was much activity in the corridor of the group studied. People walked around in the corridor in order to find specific persons, or to find someone to talk to. The results indicate that people prefer to ask a close colleague sitting next door before they contact another, more distant, person that is known to know more about the topic. This can be compared with the Answer Garden application where a question first is directed to a local expert, and if this does not yield an answer then the question is directed to a global expert (see Chapter 3). This way of asking people in an order of competence, workload etc could be an interesting design feature in a knowledge net. How to do this is an open question but workload could, e.g., be coupled to deadlines in activities. This indicates that the ordering of contacting people is related to the awareness of what people are doing.

The way the subjects walked around in the corridor can be compared to browsing. When browsing you are looking for something and you may find other interesting things along the way. To have an environment where this is possible is most useful. People need to have opportunities to make contact in a spontaneous manner.

The awareness of how available a person is for questions depends on how the group is organised in the building. In this study the “corridor office” was used with implications such as that the subjects could see other persons sitting nearby and get an awareness of what they were doing, that the subjects could see groups of people moving in the corridor on their way to a meeting or for lunch, and that the subjects could be aware of when people wanted privacy because the door was closed.

The reciprocity is socially, and not individually, “distributed”. No market of favours was observed among the subjects. There was a kind of social understanding to help

each other and being helped by others. This social understanding can be compared to the social debt that Erickson (1996) discusses. There is a slight difference between the two. A social understanding is more positive and indicates a willingness to share information. A social debt, on the other hand, has a negative sense and indicates a "must" to share information. Social debts in that sense were not detected in the group.

The availability of persons to ask was large because there did not seem to be any implicit market of favours. It was natural for the subjects to help each other with questions, with the exception of the barriers mentioned earlier. However, the author has experienced that there are organisations in which knowledge is seen as a resource that is not to be shared among "potential competitors", i.e., other colleagues. For such an organisation a knowledge net will not be usable, unless the attitude to sharing knowledge changes.

Chapter 8

Design aspects of knowledge nets

The suggested design of a knowledge net in this chapter is aimed to be used in a knowledge organisation where people usually work in projects. The three studies performed have given several ideas that is relevant for the design. Also, advantages and disadvantages with different levels of complexity of a knowledge net are discussed.

Using personal home pages as, or as the basis for, a knowledge net

The simplest way to build a knowledge net application is to let each individual have a personal home page on the Web. The Web provides an infrastructure for sharing information. As found in study I (see Chapter 5) personal home pages can be used as a simple version of a knowledge net.

The information on a personal home page can be supplied by each person—which is the case in many organisations. One feature could be to have some kind of automatically generated personal home pages with information about what project the person is involved in, written reports and contact information. Each person could then add the rest of the information by hand.

An example of such "mixed" personal home pages are used today at IPLab (the Interaction and Presentation Laboratory at Nada, KTH). Projects that a person is involved in, publications that the person has written, contact information and a picture (if available) are automatically included on a formal home page. Each project name is linked to a description, and if available the publications are linked to an on-line version. Each person can also add optional information to the page, or create a link to an informal personal home page, which is created and maintained by the person her- or himself.

These formal personal home pages, together with the informal ones, constitute a good starting point for a knowledge net. However, today the home pages at IPLab are only presented in a list sorted by person. If a person is looking for someone with a specific competence then it can be difficult to find a person who has this competence included on the home page.

Therefore, the next step would be to make the information on the personal home pages easier to search for. To make this possible some kind of search mechanism could be included on a main knowledge net page. This search mechanism should extract keywords from the personal home pages, keywords that can be used when searching for information.

The complexity of a knowledge net only consisting of personal home pages in an organisation is on the "lower level". It is easy to implement, especially if personal home pages are already used. However, facilities such as finding the "right" person to discuss certain matters with will not be present in such knowledge net. Since this is a crucial idea in the knowledge net approach a more advanced application with a higher level of complexity is recommended. However, this involves more complicated facilities to capture information.

A higher level of complexity in a knowledge net

A higher level of complexity in a knowledge net would be something like the REFERRAL WEB (Kautz et al. 1997) where not only personal home pages are used to create a network of persons, but all Web pages within the organisation are used. The REFERRAL WEB also creates a kind of relationship between persons by following links, looking for co-authors in publications, looking at organisational charts etc. Although this is a very interesting application to be used as a knowledge net it still lacks the function of being able to detect the person most suitable to answer a question.

Instead, an even higher level of complexity in a knowledge net is needed, where people can specify what they know about different topics and how much.

Entering information into the knowledge net

As much information as possible should be entered automatically in order to minimise the effort on the persons providing information. Information such as formal education, courses that have been attended, position, what project the person is working in, and what the person is working with in the project can probably give a good indication of what a person knows. There may also be a need to complement this information by hand—people may have changed field of work since graduation, they may have a specific approach to their field of work etc. This type of information can easily be scanned from personal home pages, project reports or other written documents. The topics that can be extracted from this information may have to be complemented with a description to minimise ambiguities. This must probably be done by hand, either by the person her- or himself, or by some other person.

In order for the knowledge net to give a reference to the person most suitable for the task, the information about a reference also needs some kind of rating. The rating should make it possible to detect what it is the person knows about the specific topic and how much, see Chapter 6. Ratings can always be entered by hand by each person, but if it can be done automatically then it is preferable. One way could be to scan each home page, and if the topic is found on a home page then that person gets a certain rating, an

automatic default, which should be possible to change by hand. Another alternative is to send an email including a form where ratings can be entered to the person. The form can then easily be sent back to the knowledge net. There could also be a possibility to automatically update the ratings using feedback from the users regarding how suitable the references given were. In other words, the system could learn from the users.

The ratings are important in order for the knowledge net to give references that are to the "right" persons instead of to the persons that knows most. If, for example, a rating can tell what the person can do with her or his knowledge about the topic, then the user of a knowledge net has the possibility to make more detailed searches of the competence. If the person that knows most is always referenced then this person will be overloaded with questions. Not many persons are interested in sharing one's knowledge to that extent.

Using the template "I know enough about *the topic* to *do an activity* involving *an audience* in *a setting*." used in study II (see Chapter 6) in combination with a grade on a scale is a good alternative to add some context to a reference and to measure how much a person knows about a subject. However, the study only tested the response from the person entering the information, not how the ratings could be interpreted. A question, that is still open, is how much and how good the context supplied with this method is. This is an issue that needs to be further investigated.

Another alternative is to use categorisations as used in the SPUD project (see Chapter 4). However, to categorise one's knowledge can be associated with a number of difficulties. Lansdale (1988) shows that there is a general problem in categorising items, both when to decide which categorisations that are to be used, and when to interpret the labels of the categorisations. It is difficult, if not impossible, to create category names which can be used unambiguously, and information usually falls into more than one category. He also argues that information does not fall into neat categorisation structures. Also, Malone (1983) discusses different problems when categorising information. To select a certain category was found difficult by persons he interviewed. However, he also argues that classifying information automatically, using computer facilities, simplifies the categorisation procedure. For example, multiple classification and automatic categorisation can be used. Nevertheless, to categorise peoples knowledge about a topic according to how they work with the topic, as in the SPUD project, is an interesting approach.

Much of the information can be updated automatically, e.g., changed work positions. However, some of the information must be updated by hand. To make this procedure of interest for the contributors of information it must, of course, be easy, but it must also be relevant. If the information is used by and shown to other persons, such as the information written on personal home pages, then it is reasonable to assume that people will become more motivated to update the information. This can be compared with information on personal home pages. It was found in Study I (see Chapter 5) that the subjects interviewed expected other persons to read the information on their home page, and that the subjects tended to keep their home page up-to-date.

How can entering information be rewarded?

Some people will use a knowledge net to a larger extent than they can contribute information, e.g., people who have recently been employed. There are also most likely people who will seldom use a knowledge net but who have a lot to supply, e.g., people who really are experts within their field. There will also, of course, be people who both use a knowledge net and contribute information about what they know at a "reasonable" level. It is important that there is a balance between the effort of contributing information and the benefit of using a knowledge net (cf. Grudin 1988, Grudin & Palen 1995).

Davenport & Prusak (1998) argue that there are three "natural" benefits in an organisation for "experts" to share their knowledge with others. The first is reciprocity—a person is willing to share information with others if s/he expects that the others can give information back in the future. This is supported by the findings from Study III (see Chapter 7) where it was found that there exists a kind of social understanding to help each other and to be helped. The second is repute—it is of importance for an "expert" that people within the organisation are well aware of that s/he has valuable knowledge and is willing to share what s/he knows. Having a reputation can also lead to other benefits such as a promotion. The third is altruism—many people simply have a natural impulse to help others.

If such benefits are "natural" in an organisation then they should be encouraged, with or without a knowledge net. This could be enhanced by, e.g., making it visible when people are willing to share their knowledge with others and that they are willing to do so.

Another issue, that could lower the need for reward, is that if people "own" the information about themselves, then they may feel more responsible for it. This issue can be compared with the information people write, and are responsible for, on personal home pages (see Chapter 5).

Searching for who-knows-what

The aspect of searching for who-knows-what in a knowledge net has not been an issue in any of the three studies made. The ideas presented here need to be tested in the future.

Some possibilities to search for who-knows-what about a topic in a knowledge net have been identified. One is to search in a list of topics presented in the knowledge net. Another is to have a free text search option. The first possibility is the easiest to implement, but the second may be convenient for the user. In both alternatives it should be possible to search for more than one topic. A third possibility is to use the same template as when rating the knowledge. People could enter a topic and what actions they want the person searched for to be able to perform. A fourth possibility is to use a self organising map (Honkela, Kaski, Lagus & Kohonen 1996) where statistical data are arranged so that data that are similar are in general mapped close to each other. The resulting map makes it possible to visualise distance relationships between data in an illustrative way.

Problems that can be associated with searching for who-knows-what are those of

indexing (Woods 1997). One problem is the level of depth to search. Another problem concerns the generality of topics. A topic can be associated with more than one other topic. A third problem is the name of the topic itself. There is not only one word for a topic. The same topic can have several names. These problems can lead to a large amount of possible "places" to look, i.e., one topic can be found as a subtopic of more than one other topic, depending on the number of different names of a topic it can be problematic to find it in an alphabetically ordered list, and there may be many levels of subtopics. However, some of these problems can be reduced if conceptual indexing (Woods 1997) is used. Conceptual indexing takes advantage of the conceptual structure of the information in the indexed material.

One approach used in study II (see Chapter 6) was to let the users create relations between topics. These relations can then be used when searching for information in the knowledge net. If one topic is related to another then it could be of interest to present who-knows-what about the related topic also.

The user of a knowledge net should not only be presented with one suggestion of an expert when searching for who-knows-what, but with a number of experts. People may be away and this would probably increase the time for establishing a contact with the suggested experts.

Contacting the expert

When the knowledge net has suggested a number of persons with knowledge suitable for the task, then the next step is to contact one or several of the persons. This is where the awareness information is important. An awareness of how busy the suggested experts are at the moment would ease the contacting procedure. There are many awareness cues surrounding people sharing a corridor, such as small white boards, post-it notes with messages, or people moving in the corridor (see Chapter 7). Making these awareness cues visible and integrated with the knowledge net application will be an advantage.

To make people feel more acquainted with the experts suggested by the knowledge net could also be an important issue in a knowledge net. In study I (see Chapter 5) it was found that people, if they are going to meet with a person they are unfamiliar with, often visit that person's home page in order to get a feeling of knowing the person better. It could be useful if those supplying information to the knowledge net also supply a description of themselves.

To enhance the contacting procedure there should also be messaging systems integrated with the knowledge net. If people prefer to be contacted by email then there should also be a possibility to send an email to the persons being referenced. The knowledge net could, for example, provide a possibility to directly write an email to the expert. The experts may also have preferences with respect to how they want to be contacted. Therefore, each person who has entered information about what they know into the knowledge net could also enter a kind of profile including how they prefer to be contacted.

The integration with other types of applications such as awareness, email etc are important since many systems fail because they are single applications (cf. Stewart 1997*a*).

Supporting the establishment of communities of practice

In a knowledge net, one interesting aspect is to support that people with common interests create groups where they can discuss the topic, so called communities of practice. Communities of practice is a group of co-workers who have complementary knowledge about and share an interest in a specific area (Davenport & Prusak 1998). A group of co-workers can be one way for people with similar interests in their work to meet and discuss ideas and problems. The topics entered and the ratings made on a lower level in study II (see Chapter 6) indicate that people could be interested in such a feature. If this is supported by a knowledge net then these informal networks can become organisation wide.

One way to support the establishment of communities of practice is to have a search function in the knowledge net that can find all persons with an interest about a topic. There can also be a function in the knowledge net to make a mailing list which includes all these persons.

Using the AETHER Awareness Engine in a knowledge net

The AETHER Awareness Engine (Şandor, Bogdan & Bowers 1997) is a "generic model for supporting awareness in cooperative systems". The main goal with AETHER is to recognise awareness at a fundamental system level and to build other functions on top of it. AETHER is based on the work of the Spatial Model of awareness, developed in the European Communities' COMIC project in 1992–1995 (cf. Benford & Mariani 1993, or Benford, Fahlén, Bowers & Rodden 1994). The Spatial Model has been used to enhance interaction in shared virtual environments in several experimental applications and has also influenced the architecture in a couple of Virtual Reality systems (Benford, Fahlén, Bowers & Rodden 1994). Another modification of the Spatial Model has been made by Rodden. He has generalised the Spatial Model to make it suitable for cooperative applications.

AETHER supports a semantic network of objects and directed relations between objects. The objects can be files, folders, applications, groups, people, or whatever is defined by the environment and its applications. The relations connecting different objects can be of any type defined by the environment. The semantic network of objects and relations creates a space in which an *aura*, a *nimbus* and a *focus* are defined, and in which an awareness computation is made.

The aura "describes the potential for collaboration between two objects", the focus describes how aware the object or relation is of other objects or relations, and the nimbus describes the presence of the object or relation. "Each object or relation can control its focus and nimbus to specify their 'willingness' to become aware of others or fall within their awareness". "The more an object is within your focus, the more aware you are of it. ... The more an object is within your nimbus, the more it is aware of you" (Benford et al. 1994).

The aura, focus and nimbus in AETHER also have a time component which makes it possible for the user to focus on the past, the present, or even the future. The time

component is included in what is called the medium of communication. The medium is used to "transfer" information from one object to another, as space is the medium used for transmitting radio signals from a sender to a receiver.

When computing the level of awareness one object has of another object the concept "space as an aura, nimbus and focus 'consumer' " has been added in AETHER. This means that the level of awareness between two objects depends on the nature of the space between them. This can be compared with how fog consumes part of the light and sound passing through it, and, thereby, filtering out fine details. Therefore, "The level of awareness that object A has of object B in medium M, in case of aura collision, as being some function of A's focus in M *'filtered by' the space between A and B* and B's nimbus in M *'filtered by' the space between B and A*." The emphasised text is what distinguishes the AETHER model from the Spatial Model. In other words:

> *According to this new definition, the computation of the aura, nimbus and focus becomes a negotiation between the two objects or relations (A and B), the medium that both of them 'understand' (M) and the space between the two. The medium M defines the aurae of A and B while the aura consumption is defined by the objects and relations on the relevant path(s) between them. If the aurae intersect at some point at a high enough level, then focus and nimbus computations will take place. A will define its initial nimbus value, then the different objects and relations on the path(s) between A and B will consume it. B will define the initial value of its focus and again the objects and relations on the path(s) between the place B is focusing on and A will consume it.* (Şandor et al. 1997)

This percolation of the aura from object A through the objects and relations in the medium subspace decreases the aura level each time it is consumed. The percolation continues until the aura level is below some threshold. The objects and relations that do not belong to the medium subspace are not considered in the percolation process. The nimbus percolates in the same way as the aura, while the focus on an object can be made of more than one percolation.

The critical parts of the AETHER Awareness Engine concern computing time and network size. Since no objects or relations have been deleted when they are used, because of the possibility to focus on the past, the network rapidly grows. Several ideas have been identified that deal with this problem. One idea is to use some kind of garbage collection to remove certain objects and relations. The question is which objects and relations are not significant enough to keep in the network. Another idea is to use techniques to reduce the number of computations made when searching for candidates of interest, e.g., to use parallelisation. The calculation of, for example, focus and nimbus can be done independently of each other.

Modelling a knowledge net

To build a knowledge net application using the AETHER Awareness Engine is an interesting approach. It provides awareness over a semantic net of objects and relations

that is well suited for a knowledge net application. AETHER can provide answers to questions about awareness between objects, i.e., how aware objects are of each other. This can be used in a knowledge net application to find people with knowledge about a certain topic, i.e., persons that are aware of a topic. Because of the problem with slow computation time the aim will be, in a first version, to use only a minimal model, i.e., to use only what is necessary for the purpose of the knowledge net.

In the AETHER model people and topics can be represented as objects. Each person object can have a focus on topic objects that s/he is interested in and has knowledge about. Also, each person object can have a nimbus telling how willing the person is to share her/his knowledge about the topic. However, the nimbus may not be important because the focus alone can tell that the person has knowledge about a topic. How the focus and nimbus are to be computed is still to be decided. Since no movements will be present in a knowledge net space aura will not be needed. Topic objects probably do not need any focus or nimbus.

Topic objects can be related to each other. There will be an attempt to use the relations made by the users, as in the knowledge net prototype study (see Chapter 6). Topic objects can also be related to a domain object. The purpose with domain objects is to make it possible to group topics together according to how similar they are. An example is the domain *programming languages* which could contain topic objects such as Java, C, C++ etc. The relation used will be a containment relation and it will be directed from the domain to the topic.

Person objects can, in a similar way, be related to group objects. People usually belong to a group in the organisation. The purpose with group objects is to make it possible to identify where people are located in the organisation.

The ratings given are somewhat difficult to model. One attempt is to let the grade represent a relation between the person and the topic. The activity, audience and setting can be represented as objects, but these will be generic objects. Therefore, yet another object that is concrete and that is related to the person, the topic and to the activity, audience and setting objects is needed. This will be an action object.

A model according to this is presented in Figure 8.1[1]. Person A knows about topic A and person B knows about topic B. Person A and B belong to the same group, and topic A and B belong to the same domain. Person A has defined that s/he knows enough about topic A to perform activity A involving an audience in a setting. Person B has defined that s/he knows enough about topic B to perform activity B involving the same audience in the same setting as person A's activity. Person A has rated her/his knowledge about the topic with the grade 2, and person B has rated her/his with the grade 3.

Because of the problems identified with the scale method when rating knowledge about topics in Study III (see Chapter 6) another approach will be used. There will probably be a four graded scale where 1 corresponds to "I know about this topic but I do not or have not worked with it", 2 corresponds to "I know about this topic and I am or have been working with it to some extent", 3 corresponds to "I know about this topic and I am or have been working with it on a daily basis", and 4 corresponds to "I know

[1] this work, recently started, has been done together with Cristian Bogdan

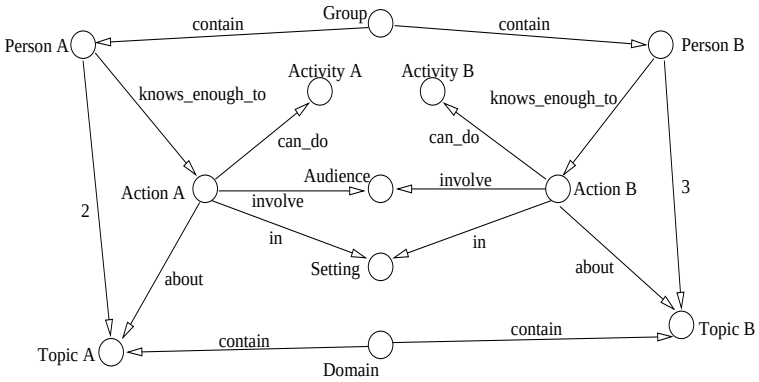

Figure 8.1: *The model of a knowledge net application to be used in the* AETHER *Awareness Engine.*

about this topic and I am considered to be an expert on it". This approach is based on parts of the rating methods used in the SPUD project (see Chapter 4).

Why then use AETHER and not an ordinary relational database? One reason is that a relational database can only supply objects and relations and not make any computations of how related the objects are. AETHER, on the other hand, can compute links at any time, either when asked for, or running in the background. The possibility to compute links in the background further makes it possible to supply users with a suggested contact net of people with knowledge about topics the user rated on a certain level. For example, if the user of a knowledge net has supplied the knowledge net with topics that s/he has knowledge of, then other persons, with similar knowledge of the same topics, can be suggested as interesting contacts. The percolation necessary to identify these contacts can run in the background, without the user being aware of it. Further, AETHER supports cooperative applications for shared or distributed systems used by multiple users. Therefore, AETHER gives a more flexible solution to a knowledge net application.

Chapter 9

Summary and conclusions

Three studies have been performed in order to collect information on how a knowledge net could be designed. Different methods have been used in the studies. Semi-structured interviews have been performed in one of the studies. In another study ethnographic methods were used, including deep interviews with open ended questions and observation of people in a real work situation. Finally, in one study people were observed, video and tape recorded while they were using a prototype of parts of a knowledge net. The diverse use of methods in the studies has been necessary in order to collect the data needed to design a knowledge net.

Interviewing three different groups of people about their personal home pages on the Web (see Chapter 5) gave important indications on what information is of interest to include in a knowledge net. Since the subjects themselves included information such as project information, contact information, and areas of interest, and also valued this kind of information on other persons' home pages, this is information that should be included in a knowledge net. The study also showed that personal home pages can be, and to some extent are, used as a simple version of a knowledge net. Another interesting result found in the study was that the subjects seemed to be quite interested in keeping their home page up-to-date. If this is because the persons themselves created and felt responsible for the content of the page then this responsibility of information would also be an interesting thing to try in a knowledge net.

Implementing a prototype of specific parts of a knowledge net and testing it among potential users (see Chapter 6) did not show any indication of the subjects not wanting to share their knowledge with others. Using the activity method when rating the knowledge seems to give a good description of people's knowledge about a topic. However, how this can be used to give a reference to the "right" person needs further investigation. The study also showed that the subjects not only rated their knowledge about topics they knew much about, but also topics that they were interested in. This indicates that some kind of support for sharing common interests would be useful to include in a knowledge net application.

Observing a group of software developers in a real work place in combination with interviews (see Chapter 7) showed that there existed socially distributed reciprocity in

the group studied. The subjects were open for questions and no situation was observed where a favour was requested in return. However, contacting people at other departments could give rise to problems. One such case was identified during the interviews—there had been a discussion on management level whether or not one person at another department was allowed to spend time on discussing a matter. Thus, sharing information within the department does not seem to be an issue in a knowledge net, but sharing knowledge over departments can be a problem. The study also identifies different awareness cues used among the persons in the corridor. Post-it notes with messages were stuck on the door and small white-boards beside the door were used for messages. The possibility to open the door to show availability or to close the door for privacy is one of the advantages with a corridor office. This suggests that the awareness of other persons is important and could also be important in a knowledge net. The study also showed that the persons studied had an ordering of whom they contacted. First they ask people nearby, within the same part of the project. If no answer is given then they ask someone further down the corridor, and so on. This indicates that the proximity is an important aspect to take under consideration in the design of a knowledge net.

One important design aspect of a knowledge net application is the level of complexity. A too low level of complexity, using only, e.g., personal home pages on the Web, would not include the functions that would facilitate finding the right person to ask. On the other hand, a too high level of complexity would render more of an organisational memory system with a rapidly growing database, with information difficult to formalise, and it would most likely be too complicated to support. What level of complexity needed depends on a number of factors such as the size of the organisation, and the complexity of the tasks carried out in the organisation, and needs further investigation. It must be complex enough to provide useful information, but not too complex to support and use.

Future work of the knowledge net approach should involve implementing the knowledge net prototype further. Functions such as searching for knowledge, retrieving relevant topics from, for example, personal home pages or project pages, and a graphical presentation of suggested experts should be tested among potential users. This may be combined with the work of using the AETHER Awareness Engine as a platform for a knowledge net application.

Bibliography

Ackerman, M. S. (1994), Augmenting the Organizational Memory: A Field Study of Answer Garden, *in* CSC (1994).

Ackerman, M. S. & Malone, T. W. (1990), Answer Garden: A Tool for Growing Organizational Memory, *in* 'Proceedings of ACM Conference on Office Information Systems', Cambridge, MA, USA, ACM Press.

Ackerman, M. S. & McDonald, D. W. (1996), Answer Garden 2: Merging Organizational Memory with Collaborative Help, *in* CSC (1996).

Anderson, J. R. (1990), *Cognitive Psychology and its Implications*, W. H. Freeman and Company. ISBN 0-7167-2085-X. 3rd edition.

Bannon, L. J. (1986), Helping Users Help Each Other, *in* 'User Centered System Design, New Perspectives on Human-Centered Interaction', Lawrence Erlbaum Associates.

Bannon, L. J. & Bødker, S. (1997), Constructing Common Information Spaces, *in* ECS (1997).

Bannon, L. J. & Kuutti, K. (1996), Shifting Perspectives on Organizational Memory: From Storage to Active Remembering, *in* 'Proceedings of HICSS-29', The Institute of Electrical and Electronics Engineers Inc (also available on CD-ROM).

Bannon, L. J. & Schmidt, K. (1991), CSCW: Four Characters in Search of a Context, *in* J. M. Bowers & S. D. Benford, eds, 'Studies in Computer Supported Cooperative Work: Theory, Practice and Design', North-Holland, pp. 3–16.

Benford, S., Fahlén, L. E., Bowers, J. & Rodden, T. (1994), A Spatial Model of Awareness, *in* 'Proceedings of the 6th ERCIM workshop', Stockholm, Sweden.

Benford, S. & Mariani, J. (1993), Requirements and Metaphors of Shared Interaction, Technical report, D4.1 COMIC Esprit Basic Research Project 6225. ftp://ftp.comp.lancs.ac.uk/pub/comic/.

Benjamins, V. R., Fensel, D. & Pérez, A. G. (1998), Knowledge Management through Ontologies, *in* Reimer (1998).

Bentley, R., Appelt, W., Busbach, U., Hinrichs, E., Kerr, D., Sikkel, K., Trevor, J. & Woetzel, G. (1997), 'Basic support for cooperative work on the World Wide Web', *Int J. Human-Computer Studies* **46**(6).

Berners-Lee, T. (1996*a*), 'Tim Berners-Lee', http://www.w3.org/pub/WWW/People/Berners-Lee-Bio.html/Overview.html (visited 5th of September, 1996).

Berners-Lee, T. (1996*b*), 'WWW: Past, Present, and Future', *Computer* **29**(10), 69–77. http://www.computer.org/computer/co1996/rx069abs.htm.

Berners-Lee, T., Cailliau, R., Luotonen, A., Frystyk Nielsen, H. & Secret, A. (1994), 'The World-Wide Web', *Communications of the ACM* **37**(8).

Bly, S., Cook, L., Bickmore, T., Churchill, E. & Sullivan, J. W. (1998), The Rise of Personal Web Pages at Work, *in* 'Summary of CHI'98', ACM Press.

Bogdan, C. & Şandor, O. (1997), JML—The New CGI Programming Solution. http://best.nada.kth.se:8080/jml/jml.cgi/index.jml (visited May 20th, 1998).

Bowers, J. (1994), The Work to Make a Network Work: Studying CSCW in Action, *in* CSC (1994).

Bowers, J. (1995), Hanging around and making something out of it, *in* J. T. Haworth, ed., 'Psychological Research: Innovative Methods and Strategies', Routledge, UK, chapter 8.

Bowker, G. C. (1996), Lest We Remember: Organizational Forgetting and the Production of Knowledge. http://alexia.lis.uiuc.edu/~bowker/forget.html (visited September 1998).

Clark, H. (1996), *Using Language*, Cambridge University Press.

Collins, H. M. (1985), *Changing Order*, SAGE Publications.

Conklin, E. J. (1996), Designing Organizational Memory: Preserving Intellectual Assets in a Knowledge Economy. http://aroi.ing.puc.cl/~dfuller/conklin.html (visited May 1998).

Conklin, J. & Begeman, M. L. (1988), gIBIS: A Hypertext Tool for Exploratory Policy Discussion, *in* CSC (1988).

Coulter, J. (1989), *Mind in Action*, Polity Press.

Crossley, M., Davies, J., McGrath, A. & Rejman-Greene, M. (1998), The Knowledge Garden, *in* Reimer (1998).

Şandor, O., Bogdan, C. & Bowers, J. (1997), Aether: An Awareness Engine for CSCW, *in* ECS (1997).

CSC (1988), *Proceedings of CSCW'88*, ACM Press.

CSC (1994), *Proceedings of CSCW'94*, ACM Press.

CSC (1996), *Proceedings of CSCW'96*, ACM Press.

Davenport, T. H. & Prusak, L. (1998), *Working Knowledge*, Harvard Business School Press. ISBN 0-87584-655-6.

Degerblad, J.-E. (1988), Yrkeskunskaper under svenskt 1700-tal, *in* B. Göranzon, ed., 'Den inre bilden', Carlssons Bokförlag.

Densford, L. E. (1996), 'Knowledge Management; an Underutilized Tool That Enables Companies to Know What They Know', *Corporate University Review* (December).

Dix, A. (1996), Challenges and Perspectives for Cooperative Work on the Web, *in* 'Proceedings of the ERCIM workshop on CSCW and the Web', Sankt Augustin, Germany, GMD/FIT.

Dix, A., Finlay, J., Abowd, G. & Beale, R. (1998), *Human-Computer Interaction*, Prentice Hall Europe. ISBN 0-13-239864-8. 2nd edition.

Dourish, P. (1993), Culture and Control in a Media Space, *in* 'Proceedings of ECSCW'93', Kluwer Academic Press.

Dourish, P. & Belotti, V. (1992), Awareness and Coordination in Shared Workspaces, *in* 'Proceedings of CSCW'92', ACM Press.

ECS (1997), *Proceedings of ECSCW'97*, Kluwer Academic Publishers.

Ellis, C. A., Gibbs, S. J. & Rein, G. L. (1991), 'Groupware: Some Issues and Experiences', *Communications of the ACM* **34**(1).

Erickson, T. (1996), 'The World-Wide Web as Social Hypertext', *Communications of the ACM* **vol. 39**(1), 15–17.

Eveland, J. D., Blanchard, A., Brown, W. & Mattocks, J. (1994), The Role of "Help Networks" in Facilitating Use of CSCW Tools, *in* CSC (1994).

Ford, K. M., Bradshaw, J. M., Adams-Webber, J. R. & Agnew, N. M. (1993), 'Knowledge Acquisition as a Constructive Modeling Activity', *International Journal of Intelligent Systems* **8**, 9–32.

Gantt, M. & Nardi, B. A. (1992), Gardeners and Gurus: Patterns of Cooperation Among CAD Users, *in* 'Proceedings of CHI'92', ACM Press.

Groth, K. (1997), The Use of Knowledge Nets for Collaboration within Organizations: a Theoretical Background, Technical Report IPLab-127, TRITA-NA-P9703, Royal Institute of Technology, Sweden.

Groth, K. (1998), 'Personal Home Pages on the World Wide Web — a Simple Version of a Knowledge Net?', *Trends in Communication* (4), 47–59.

Grudin, J. (1988), Why CSCW Applications Fail: Problems in the Design and Evaluation of Organizational Interfaces, *in* CSC (1988).

Grudin, J. & Palen, L. (1995), Why Groupware Succeeds: Discretion or Mandate?, *in* 'Proceedings of ECSCW'95', Stockholm, Sweden, Kluwer Academic Publishers.

Gundry, J. & Metes, G. (1996), A Working by Wire. http://www.knowab.co.uk/wbwteam.html.

Hollocks, B. W. (1998), The Role of Models in Leveraging Information, *in* Reimer (1998).

Honkela, T., Kaski, S., Lagus, K. & Kohonen, T. (1996), Newsgroup Exploration with WEBSOME Method and Browsing Interface, Technical Report A32, Helsinki University of Technology, Faculty of Information Technology. ISBN 951-22-2949-8.

Hughes, J., Rodden, T. & Rouncefield, M. (1994), CSCW Requirements Development, Technical report, D2.2 COMIC Esprit Basic Research Project 6225. ftp://ftp.comp.lancs.ac.uk/pub/comic/D2.2.ps.Z.

Instone, K. (1996), 'HCI and the Web', *SIGCHI Bulletin* **28**(4).

Kautz, H., Selman, B. & Shah, M. (1997), 'The Hidden Web', *AI Magazine, AAAI* pp. 27–35.

Kjellin, H. (1994), A Method for Acquiring and Refining Knowledge in Weak Theory Domains, PhD thesis, Dept of Computer Systems, Stockholm University, Sweden.

Kraut, R. E., Egido, C. & Galegher, J. (1988), Patterns of Contact and Communication in Scientific Research Collaboration, *in* CSC (1988).

Kraut, R. E., Galegher, J. & Egido, C. (1987–1988), 'Relationships and Tasks in Scientific Research Collaboration', *Human-Computer Interaction* **3**, 31–58.

Kraut, R. E. & Streeter, L. A. (1990), Satisfying the Need to Know: Interpersonal Information Access, *in* 'Proceedings of Interact'90', Elsevier Science Publishers.

Kraut, R. E. & Streeter, L. A. (1995), 'Coordination in Software Development', *Communications of the ACM* **vol. 38**(3), 69–81.

Kühn, O. & Abecker, A. (1997), 'Corporate Memories for Knowledge Management in Industrial Practice: Prospects and Challenges', *Journal of Universal Computer Science* **3**(8), 929–954. Springer Pub. Co.

Lansdale, M. (1988), 'The Psychology of Personal Information Management', *Applied Ergonomics* **19**(1), 55–66. Butterworth & Co Ltd.

Lochbaum, K. E. & Streeter, L. A. (1989), 'Comparing and combining the effectiveness of latent semantic indexing and the ordinary vector space model for information retrieval', *Information Processing & Management* **25**(6), 665–676. Pergamon Press plc.

Malone, T. W. (1983), 'How Do People Organize Their Desks? Implications for the Design of Office Information Systems', *ACM Transactions on Office Information Systems* **1**(1), 99–112.

Marmolin, H. (1991), TheKnowledgeNet, *in* 'Proceedings of the 2nd MultiG Workshop', Royal Institute of Technology, Stockholm, Sweden.

Marmolin, H. & Sundblad, Y. (1993), Sharing Knowledge in a Distributed Environment for Collaboration, *in* 'Proceedings of the 6th MultiG Workshop'.

Marmolin, H., Sundblad, Y., Lidbaum, P. & Tobiasson, M. (1991), The Collaborative Desktop: Interface and Functionalities, *in* 'Proceedings of the 3rd MultiG Workshop'.

Marmolin, H., Sundblad, Y., Tollmar, K., Avatare, A. & Eriksson, H. (1992), CoDesk — an Interface to TheKnowledgeNet, Design and Implementation, *in* 'Proceedings of the 4th MultiG Workshop'.

McDonald, D. W. & Ackerman, M. S. (1998), Just Talk to Me: A Field Study of Expertise Location, *in* 'Proceedings of CSCW'98', ACM Press.

Polanyi, M. (1964), *Personal Knowledge*, Harper & Row Publishers Inc.

Pycock, J. & Bowers, J. (1996), Getting Others To Get It Right: An Ethnography of Design Work in the Fashion Industry, *in* CSC (1996).

Reimer, U., ed. (1998), *PAKM, Practical Aspects of Knowledge Management*.

Rittel, H. & Kunz, W. (1970), Issues as elements of information systems, Technical report, Working paper #131. Institut für Grundlagen der Planung I.A., University of Stuttgart.

Rodden, T. (1997), Populating the Application: A Model of Awareness for Cooperative Applications, *in* ECS (1997).

Ryle, G. (1949), *Concept of Mind*, William Brendon and Son, Ltd.

Salvador, T. & Bly, S. (1997), Supporting the Flow of Information Through Constellations of Interaction, *in* ECS (1997).

Stewart, T. A. (1997*a*), Does Anyone Around Here Know ... ? http://www.pathfinder.com/fortune/1997/970929/lea.html (visited February 1998).

Stewart, T. A. (1997*b*), *Intellectual Capital: The New Wealth of Organizations*, Doubleday, New York. ISBN 0-385-48228-0.

Streeter, L. A. & Lochbaum, K. E. (1988), An Expert/Expert-Locating System Based on Automatic Representation of Semantic Structure, *in* 'The Fourth Conference on Artificial Intelligence Applications', IEEE.

Tauscher, L. & Greenberg, S. (1997), 'How People Revisit Web Pages: Empirical Findings and Implications for the Design of History Systems', *Int. J. Human-Computer Studies* **vol 47**, 97–137.

Terveen, L. G., Selfridge, P. G. & Long, M. D. (1993), From "Folklore" to "Living Design Memory", *in* 'Proceedings of INTERCHI'93', ACM Press.

Tollmar, K., Marmolin, H. & Sundblad, Y. (1994), The Collaborative Desktop—An Environment for Computer Supported Cooperative Work, *in* 'Conference Companion of CHI'94', ACM Press.

Tollmar, K. & Sundblad, Y. (1995), 'The Design and Building of the Graphic User Interface for the Collaborative Desktop', *Computers and Graphics* **vol. 19**(2), 179–188.

Trevor, J., Koch, T. & Woetzel, G. (1997), MetaWeb: Bringing synchronous groupware to the World Wide Web, *in* ECS (1997).

Wærn, Y. & Wærn, K.-G. (1984), *Tänkande pågår — om tankepsykologi för ADB-folk*, Liber Förlag. ISBN 91-38-61297-6.

Walsh, J. P. & Ungson, G. R. (1991), 'Organizational Memory', *Academy of Management Review* **16**(1), 57–91.

Waterson, P. E., Clegg, C. W. & Axtell, C. M. (1997), 'The dynamics of work organization, knowledge and technology during software development', *Int. J. Human-Computer Studies* **vol. 46**, 79–101.

Woods, W. A. (1997), Conceptual Indexing: A Better Way to Organize Knowledge, Technical report, Sun Microsystems Laboratories.

Appendix A

Personal home pages on the Web

Interview questions 1996

1. Background on how the Web is used
    - When did you start to use the Web?
    - Do you use the Internet in any other way? If so, how?
    - How often do you use the Web?
    - For how long time do you usually use the Web?
    - For what do you use the Web (look for information, develop, surfing)?
    - If searching for information, what kind of information do you look for?
    - If searching for information, do you often look for the same kind of information?
    - If searching for information, do you plan ahead what it is you are looking for?
2. Background of the home page
    - When did you write your home page?
    - Why did you write a home page?
    - How many times have you made major changes to your home page?
    - What kind of changes did you do (links, pictures, text, layout)?
    - Why did you do these changes?
3. Are people satisfied with their home page?
    - Are you, on the whole, satisfied with your home page?
    - What are you (not) satisfied with?

- Are you satisfied with the layout of your home page?
- What are you (not) satisfied with?
- Are you satisfied with the content on your home page?
- What are you (not) satisfied with?
- If you have any links on your home page, are they for you or for visitors?

4. Are there any limitations on who can read the home page?

   - What language have you used on your home page (Swedish or English)?
   - Why did you choose that language?

5. Have any other persons home pages inspired to the home page?

   - What influenced the layout of your home page (an organisation template, own design, other persons' home page, a mixture of other persons' home pages)?
   - What influenced the content of your home page (an organisation template, own design, other persons' home pages, a mixture of other persons' home pages)?
   - What influenced the choice of language on your home page (an organisation template, own design, other persons' home pages, a mixture of other persons' home pages)?
   - Did you look at other persons home pages before you made your own? If so, can you give any examples?
   - Where were these home pages (inside or outside your own organisation)?
   - Did you use another person's home page as a template when you made your own? If so, give some examples.
   - Why did you use another person's home page as a template?

6. What persons are seen as potential readers of the home page?

   - What persons or groups of persons do you think can be interested in reading your home page?
   - Why do you think these persons were interested in reading your home page?
   - What is the purpose with your home page, what is the message?

7. What is found on the home page?

   - What kind of information have you included on your home page (contact information, projects that I am working with, areas of interest within work, hobbies, home address, picture of myself, picture of my family)?

- Why did you (not) include contact information, projects that I am working with, areas of interest within work, hobbies, home address, picture of myself, picture of my family?

8. What is found to be of interest on a personal home page?

    - How many times have you been contacted by persons who have seen your home page?
    - What information seems to be of most interest among persons contacting you?
    - Do you find information on other persons' home pages useful?
    - What information do you find interesting on other persons' home pages?

Interview questions 1997

These interviews were conducted as a discussion around the old and the new home page. However, questions that should have been answered at the end of the interview were:

1. What projects are you involved in right now?

    - Are these projects described on any other page than your own home page?
    - If yes: Do you have a link to this description from your home page? How come?
    - If no: Do you describe these projects on your home page? How come?
    - How do you or would you feel about presenting this kind of information on the Internet?
    - How do you or would you feel about presenting this kind of information on an intranet?
    - Would it make any difference if everybody else within the organisation presented this kind of information?

2. What projects have you been involved in during your time at this workplace?

    - Are these projects described on any other page than your own home page?
    - If yes: Do you have a link to this description from your home page? How come?
    - If no: Do you describe these projects on your home page? How come?
    - How do you or would you feel when presenting this kind of information on the Internet?
    - How do you or would you feel when presenting this kind of information on an intranet?

- Would it make any difference if everybody else within the organisation presented this kind of information?

3. What do you write about your educational background? Why?

   - How do you or would you feel about presenting this kind of information on the Internet?
   - How do you or would you feel about presenting this kind of information in an intranet?
   - Would it make any difference if everybody else within the organisation presented this kind of information?

4. Are you involved in other activities that can be of interest for other persons (within and without the organisation)? What kind of activities?

   - Are these activities described on any other page than your own home page?
   - If yes: Do you have a link to this description from your home page? How come?
   - If no: Do you describe these projects on your home page? How come?
   - How do you or would you feel about presenting this kind of information on the Internet?
   - How do you or would you feel about presenting this kind of information on an intranet?
   - Would it make any difference if everybody else within the organisation presented this kind of information?

5. Have you been contacted by a person that had found something on your home page that you had written for that category of persons? What kind of information was this?

6. Last time you said xxx was information that you found interesting on other personal home pages. What exactly are you interested in and why?

Appendix B

Analysis of the results from study II

t-tests

t-tests have been performed on the results from study II (see Chapter 6) in order to see if the differences between the two groups are statistically significant. The variables $n1$, $M1$ and $S1$ contain the values of the group of researchers, while the variables $n2$, $M2$ and $S2$ contain the values of the group of software developers.

Although *t*-tests have been performed and the two groups are compared, it should be noted that this comparison was not planned when the study was started. Because of the changes in the prototype interface and time to perform a session a comparison between the groups is difficult to do.

t-tests performed on the results of entering topics show that there is a statistically significant difference at the 95% level between the two groups in the amount of topics entered in session 2, see Table B.1. The difference can be explained by the fact that the software developers used the prototype in the second session for half an hour while the researchers used it for 45 minutes. The difference of the amount of topics entered in the first session was not statistically significant at the 90% level. Both groups used the prototype for 30 minutes in session 1.

Topics	$n1$	$n2$	$M1$	$M2$	$S1$	$S2$	fg	Sp	t
Session 1	16	11	5.38	6.09	2.31	2.43	25	5.55	-0.33
Session 2	14	10	1.14	0.30	1.10	0.79	22	0.97	2.10

Table B.1: *t-tests performed on the number of topics entered in both sessions by the two groups of subjects. Only the difference in session 2 was statistically significant at the 95% level.*

t-tests have also been performed on the amount of ratings made, see Table B.2, on the kind of ratings made using the scale method, see Table B.3, on the number

of relations made, see Table B.4, and on the number of topics related, see Table B.5. None of these tests showed a statistical significance at the 90% level. It is interesting though that there were no significant difference regarding the amount of ratings entered in session 2. This could mean that the subjects turned to entering topics again when they had more time.

Topics	$n1$	$n2$	$M1$	$M2$	$S1$	$S2$	fg	Sp	t
Session 1	16	11	5.06	5.36	2.26	2.87	25	6.38	-0.12
Session 2	14	10	4.79	3.70	3.81	2.91	22	12.02	0.22

Table B.2: *t-tests performed on the number of ratings entered in both sessions by the two groups of subjects. These differences were not statistically significant at the 90% level.*

Topics	$n1$	$n2$	$M1$	$M2$	$S1$	$S2$	fg	Sp	t
Session 1	16	11	5.12	4.73	0.48	1.17	25	0.69	1.43
Session 2	14	10	4.92	4.53	0.77	1.76	22	1.62	0.58

Table B.3: *t-tests performed on the kind of ratings made using the scale method in both sessions by the two groups of subjects. These differences were not statistically significant at the 90% level.*

Topics	$n1$	$n2$	$M1$	$M2$	$S1$	$S2$	fg	Sp	t
Total	14	10	34.71	11.1	37.68	9.27	22	874.27	0.07

Table B.4: *t-tests performed on the number of relations entered by the two groups of subjects. This difference was not statistically significant at the 90% level.*

Topics	$n1$	$n2$	$M1$	$M2$	$S1$	$S2$	fg	Sp	t
Total	14	10	17.6	5.6	13.7	4.1	22	118.33	0.24

Table B.5: *t-tests performed on the number of topics related by the two groups of subjects. This difference was not statistically significant at the 90% level.*